University of Alberta

Library Release Form

Name of Author: James Andrew Smith

Title of Thesis: A Tracking System for a Seismic Surveying Mobile Robot

Degree: Master of Science

Year this Degree Granted: 2001

"I treated the whole field broadly, not limiting myself to mechanisms controlled from a distance but to machines possessed of their own intelligence. Since that time I had advanced greatly in the evolution of the invention and think that the time is not distant when I shall show an automaton which, left to itself, will act as though possessed of reason without any willful control from outside. Whatever be the practical possibilities of such an achievement, it will mark the beginning of a new epoch in mechanics." [1]

Nikola Tesla (1856 - 1943) describing his pioneering work in robotics in the 1890's.

University of Alberta

A Tracking System for a Seismic Surveying Mobile Robot

by

James Andrew Smith ©

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of **Master of Science**.

Department of Electrical and Computer Engineering

Edmonton, Alberta
Spring 2001

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled **An Outdoor Mobile Robot for the Seismic Surveying Industry** submitted by James Andrew Smith in partial fulfillment of the requirements for the degree of **Master of Science**.

Cette thèse est dédier à mes parents.

Abstract

This thesis presents the synergistic application of visual and SONAR sensing to a full-size robotic land vehicle, enabling the vehicle to follow a cooperating human leader, on foot, in an unstructured, outdoor environment. The robot determines the range and bearing to the leader by combining the output of the visual sensors, which track the leader's distinctly marked target vest, and the ultrasonic sensors, which receive signals from an ultrasonic beacon worn by the leader. Unlike similar "follow-the-leader" robotic projects, fixed latency radio synchronization between the robot and the leader's ultrasonic beacon is used, providing more accurate range data. This system also distinguishes itself in its application of a set of wide field-of-view SONAR transducers, permitting less constrained operation than those employing acoustic reflectors. These two factors were critical, as the ultimate goal of the project was to demonstrate a reliable and inexpensive approach to semi-autonomous ("follow-the-leader") control applied to a full-sized (400 kg), hydraulically powered, robotic land vehicle equipped with an industrial payload, in moderately complex, unstructured outdoor environment, at speeds of up to 5 km/h and with a minimum turning radius of 3 m. In conclusion, the concept was successfully demonstrated, and several avenues for further research were identified.

Preface

Choosing the right tools and learning how to use them properly will almost always result in a better product: in this case, a thesis. Often, these tools are not brand-new, with all the latest bells and whistles and commercial product support. Writing a thesis is a difficult enough process without having to contend with bugs in the latest proprietary software. The following is a list of some of the tools I used to write my thesis:

- Emacs, a wonderful everything-but-the-kitchen-sink text editor
- LaTeX, the typesetting language used to process this document
- Xfig, the vector-based graphics development program
- The Gimp, an Adobe Photoshop clone for bitmap-based graphics development
- "How to Organize Your Thesis" by Carleton University's Prof. Chinneck *http://www.sce.carleton.ca/faculty/chinneck/thesis.html*
- The University of Alberta Thesis Document Style for LaTeX by Fahiem Bacchus et alii

The Emacs text editor is almost as old as I am and is under constant development by a talented band of programmers. This means that it is rock-solid and is filled with many useful features, including LaTeX compatibility.

LaTeX is responsible for document layout, allowing the writer to concentrate on writing. I am a graduate student in electrical engineering, not typesetting; as such, I trust the judgement of the typesetting experts who wrote LaTeX, just as I would expect someone to trust my judgement in matters related to electrical engineering.

The LaTeX learning curve is relatively steep but the time put in to learn it at the beginning of your thesis writing adventure will pay off in the end. You won't have to triple-check the accuracy of page numbering in your table of contents, or the numbering of your figures and equations. They will be correct and properly formatted. LaTeX's layout of equations is also second to none. I used the LaTeX thesis template endorsed by the Faculty of Graduate Studies, which helped speed up the writing of my thesis.

Dr. Chinneck's thesis organisation webpage receives about 150 hits per day. The reason for its popularity is simple: graduate students are often left to fend for themselves when it comes to writing theses.

Acknowledgements

This thesis would not have been possible without the help of a great number of people. I would like to take a bit of time to thank these people.

This research was funded in large part by a Defence Industrial Research Program grant. Thank you to Schlumberger Geco-Prakla, the Natural Sciences and Engineering Research Council and the Department of National Defence for supporting this research. I would also like to thank the many corporate sponsors who supported the Autonomous Robotic Vehicle Project, including MacDonald Dettwiler Robotics, Corel, Atco, and Syncrude.

Dr. Toogood, thank you for all the advice. I'm not the only one who really appreciated your version of Hans Christian Andersen's "The Emperor's New Clothes."

To my partners in crime on the SAGV project, Aaron Saunders, Ryan Chladny, Jim Qualie and Kit Barton: thanks for opening my eyes, questioning my methods, and supporting me when the magic blue smoke escaped from my circuits.

Jason Gunthorpe, thanks for the cookies! Jeff Woo, thank you for taking over the ARVP. To the organisers of the AUVSI Intelligent Ground Vehicle Competition: thank you for putting such an awesome competition together.

Thank you to the Leonard Swanson, Lynn Vogelesang and Jill Stanton in the Engineering Dean's Office. It was a real pleasure working with you during my stay in Alberta. Also, thank you to Michelle Lock, Carla Zittlau, Carla DeJager and Maureen Lebrecht in the Electrical and Computer Engineering Department office.

Loren Wyard-Scott, thank you for all the help, for the chit-chat, for not killing me when I blew up your robot, and for letting me crash at your place!

Michelle Huth, the most wonderful woman in my life, thanks for putting up with me! I'm so glad that we got to spend time in Guadalajara and Guanajuato, stomping on the cookarachas, discovering the egg-throwing and arm-flapping locals, and eating fresh potato chips in the afternoons. Pass the lime and hot sauce!

Thank you to my parents, Don and Thérèse, for the financial and moral support that helped me get through my undergrad, and then supporting me when I decided to continue on with grad studies.

Thomas, my soon-to-be-married little brother, thanks for being a real inspiration in my becoming an electrical engineer!

Finally, I would also like to separately thank my two proof-readers, Michelle Huth and my father, Donald Smith. They didn't have to wade through my thesis and I really appreciated their help in spotting spelling, grammatical and punctuation errors. They were invaluable in helping me focus on the important technical and structural issues that one must face when writing an engineering thesis.

Contents

List of Figures

Chapter 1

Introduction

1.1 The Requirement

In late 1998 the author was contacted by Allan Chatenay, Canada Land Manager for Schlumberger Geco-Prakla. As a member of the Oilfield Services business segment of Schlumberger, Geco-Prakla conducts seismic surveying work in the oilfield. Mr. Chatenay had read an article about the University of Alberta's Autonomous Robotic Vehicle Project and was struck by the idea of employing robotic vehicles in the oilfield.

Mr. Chatenay felt that by automating certain geophysical tasks Geco-Prakla could improve the delivery of its surveying data product. The obvious result of this would be that the company would be able to obtain a larger share of the geophysical surveying market.

Geophysical work is often physically demanding and labour intensive. Although the seismic surveying segment of the geophysical industry has adopted labour-saving information technologies for data processing it has been slower to implement automation for many of the more physically demanding tasks. The seismic surveying industry could benefit from targeted automation of certain laborious tasks, in particular the transportation, deployment and collection of seismic cable. This thesis addresses the development of a set of tools to enable the automation of these tasks.

Augment the Field Worker

No amount of research and development will completely eliminate human beings from work in the oilfield. The unstructured, dynamic and often unpredictable nature of field work requires the use of competent and quick-thinking people. The goal of using robots in the

oilfield is to improve the effectiveness of people already there. By targeting specific labour intensive or mundane tasks for automation it will be possible to increase the effectiveness of field workers.

Field worker augmentation is already occurring. The Navpac (as discussed on page 84), equipped with global positioning and an inertial navigation subsystems, is helping Schlumberger Geco-Prakla deliver a better surveying data product to its customers. The introduction of automation in other aspects of seismic surveying is a natural extension of what is already occurring with the Navpac.

NSERC & DND: Important Partners

In 1999 the author submitted a successful request for a Defence Industrial Research Program grant, thereby involving both the Natural Sciences and Engineering Research Council (NSERC) and Department of National Defence (DND) in the Schlumberger research project. Both of these organisations feel that there is potential for gains in the commercial, research and defence sectors through this research.

1.2 The Approach

The platform on which the research is conducted is the Polar Bear, a 400 kg. (900 lb.), gasoline powered and hydraulically actuated robotic vehicle. Development of the Polar Bear robot began in 1997 and will continue until at least the summer of 2001. It is a practical system for conducting land-based outdoor mobile semi-autonomous robotics research such as that required by Schlumberger Geco-Prakla.

Using the Polar Bear as a demonstration platform, the goal of this project is to develop a prototype pursuit robot equipped with a SONAR array and a video camera for sensory feedback. The SONAR returns range information about obstacles and the camera gives it visual feedback.

1.2.1 The Field Worker Pursuit Problem

The Polar Bear is being developed to be used by oilfield workers in a semi-autonomous context. A field worker will be assigned the task of passively leading the robot as it conducts a task such as the deployment of seismic cable.

The semi-autonomous mode of operation augments the robot with the higher-level navigational reasoning of the human leader. Additionally, the robot augments the person by performing the more labour-intensive equipment transportation or deployment task.

The link between the human operator and the robot is achieved through dual-sensory identification. The robot determines the direction to the person by identifying the person's safety jacket via the on-board camera. Range to the person is determined using radio-synchronised SONAR receivers on a fanny-pack worn by the person. The SONAR system is a unique combination of devices that is especially suited for the particular application.

Although this thesis describes the robot platform, as well as the type of sensors and manner in which they are used, it is the latter two aspects which are the main focus. The platform and the environment in which it operates provide the context for the sensors.

Because of the parties involved in this project, much emphasis has been placed on practical implementations for the sensors and platform. The demonstrations of the system required that the robot be able to follow a person outdoors relatively reliably. Therefore, the sensors and other equipment used had to perform reliably while providing necessary performance.

1.2.2 The Present State of the Project

Currently, the project is at Level 6 ("System/subsystem model or prototype demonstration in a relevant environment") out of 9 on the NASA Technology Readiness Levels scale. As was seen in a demonstration on September 25 2000, the Polar Bear robot is presently capable of following a person using SONAR and vision sensors in an outdoor environment. Speeds of up to 5 km/h on terrain containing traversable obstacles such as roadside curbs and sparse non-traversable obstacles such as trees have been attained.

1.3 State of the Art: Vision and SONAR Sensors for Robots

When considering a robotics project the general problem is constrained in three ways:

- Mobility: fixed or mobile,
- Environment: land, air, sea or space,

- Level of Autonomy: tele-operated, semi-autonomous or fully autonomous.

This project focuses primarily on mobile, land-based and semi-autonomous robotic systems. Therefore, the sensors used must be able to operate effectively under these constraints. The sensors must be relatively light weight, rugged and also low power. In addition, they should be relatively inexpensive, in order to meet the project's budgetary requirements.

Visual and ultrasonic sensors are widely available and, because many models are used in consumer systems, are relatively inexpensive and robust. Visual sensors are often used to obtain two-dimensional information from the environment in a passive manner. Through the use of multiple cameras or by making assumptions about the size and shape of objects detected by a single camera, it is possible to obtain three-dimensional information, as well.

Ultrasonic sensors, which most atmospheric SONAR systems use, can be used to determine distance to an object, as well as object size, shape, and even material composition. Multiple ultrasonic SONAR sensors can be used to determine three-dimensional location information.

In this project two sets of ultrasonic SONAR sensors are used to determine distances to objects. A camera is used in conjunction with the SONAR to determine direction. A precursor to that project uses a single camera to determine the three-dimensional distance and pose to a known target. An overview of projects related to these applications is discussed below.

1.3.1 Vision

There are two reliable ways to obtain three-dimensional information using passive video: two or more cameras arranged in such a way that the difference in images they capture can yield depth information, or to use one camera to extract visual feature point information from a target of known dimensions. In the latter case the popular Perspective-n-Point algorithm is often used.

With a minimum of three known feature points on a known coplanar target it becomes possible to use the Perspective-n-Point (n is the number of feature points on the target) pose estimation algorithm [2] and [3, errata] to determine the pose of the target. The P-n-P algorithm yields the straight-line distance to feature points on the target. The easiest P-n-P

algorithm to implement is the Perspective-3-Point (i.e. the target has three feature points) which yields up to eight possible solutions for a target's pose. Attempts have been made [4] to develop P-4-P (four feature points) and P-5-P (five feature points) using similar but somewhat more involved calculations. The validity of the work in [4] could not be easily verified since the paper contained numerous errors.

1.3.2 SONAR in General

Sound navigation and ranging (SONAR) technology has been available since about 1918. Most applications of SONAR systems measure the time taken by a pulse of sound to travel from a transmitter to a receiver. Often, this sound reflects off an object, allowing the distance to the object to be measured. Most SONAR research has been conducted on underwater implementations (e.g. submarine navigation and the detection of fish for commercial fish operations) but this thesis uses SONAR under atmospheric conditions where the sound travels more slowly and the coupling between the SONAR transducer and air, the medium in which the sound travels, is less efficient. Unlike underwater SONAR systems, atmospheric SONAR can be used in conjunction with high-frequency radio systems, as will be shown in Section 1.3.4, allowing radio synchronisation to occur between SONAR receiver and transmitter.

SONAR[1] lends itself well to the basic tasks of range-finding and target recognition. Polaroid Corporation has developed a series of SONAR transducers that are inexpensive, reliable and widely used. Although they have a wide variety of transducers available, the most commonly used in commercial applications (e.g. auto-focus cameras, door openers and mobile robots) are Polaroid's low-cost 600 series electrostatic transducers driven by the 6500 series ranging boards. The 7000 series electrostatic and 9000 series piezoelectric transducers are also available and can also be driven by the 6500 ranging boards.

It seems that most researchers and end-users use the SONAR modules for range-finding under relatively benign atmospheric conditions. In these cases the electrostatic 600 series or 7000 series transducers are used. The 9000 series transducer is rarely used outside the automotive industry, where it is mounted on car and truck bumpers for low speed obstacle detection. Although piezoelectric transducers are sometimes used in robotics applications

[1] Unless otherwise noted, the term SONAR refers to atmospheric SONAR systems.

none have been found to use Polaroid's 9000 series piezoelectric transducer.

Piezoelectric transducers couple well to solids and liquids but poorly to air, thus reducing the strength of the signal and maximum detectable distance in a range-finding application. Combined with a possible increase in latency due to the mechanical inertia of the piezoelectric crystal these types of transducers are often not as desirable as electrostatic transducers in range-finding applications.

A discussion with an engineer from iRobot Corporation in July 2000 yielded some interesting details regarding the 9000 series transducer. The company had attempted to use the more environmentally robust 9000 series but apparently, because of the properties of the transducer they had failed to field the unit in a robot. Probably due to the low strength of the transmitted SONAR pulse, only a select number of materials including some metals but not wood, gyprock, or cardboard could be detected.

SONAR sensors can also be used to send information between two different systems. One such application, beaconing, allows one system to determine the distance to another by measuring the time taken by a sound pulse (propagating at a fixed and known speed) to travel between them. Beaconing can be taken further, allowing three or more receivers to determine the three-dimensional position of the transmitter.

1.3.3 SONAR Beaconing

> "Although special beacons are at odds with notions of complete robot autonomy in an unstructured environment, they offer advantages of accuracy, simplicity, and speed - factors of interest in industrial and office applications, where the environment can be partially structured." [5, quote by Lindsay Kleeman on page 151]

When attempting to localise an object using SONAR, triangulation is the method most often used. By strict definition to triangulate is to: "measure, map out, by measurement of sides and angles of series of triangles on determined base-line(s)" [6]. In practice, what is strictly known as triangulation is broken down into two methods: triangulation and trilateration. Both use triangular relationships to determine position.

Triangulation [7] is, in practice, thought of as the determination of position based on angular measurements to at least three known locations. It is this method that surveyors or topographers use when employing a compass, a map and other sighting instruments.

Trilateration [8] uses range measurements to determine position, rather than directly measuring angular ones. Again, triangular relationships are required and so position is determined using range measurements to at least three known points. Global positioning systems (GPS) use trilateration to determine the position of GPS receivers by measuring the distance between the receiver and multiple transmitting satellites. Because range determination is so easy with SONAR, such sensors can be used in trilateration applications as well.

1.3.4 SONAR Beacon Example 1: Millibots

A team of Carnegie Mellon University researchers are working on an ongoing distributed mobile robotic project called "Millibots" [9]. These palm-sized robots are meant to be used in applications not suitable to large- or medium-sized robots.

The individual robots share sensor information as it is gathered. In order to properly interpret the combined information the researchers feel that it is important that the relative position between the robots or their position in a global coordinate system correlate with sensor data.

The Millibot localisation system works by using SONAR range measurements between robots of unknown position and robots of known position. SONAR transducers are mounted vertically beneath acoustic reflectors on the top level of each robot. These beaconing units can operate as either transmitters or as receivers and can be seen in Figure 1.1 on page 8.

1.3.5 SONAR Beacon Example 2: Ghengis

Used as a reference by the Millibot project researchers, IS Robotics (now iRobot Corporation) developed a two-dimensional location system using SONAR trilateration and adapted it to their Ghengis[2] robot. According to [5] the project's principle sources of error included "... variations in the speed of sound, the finite size of the ultrasonic transducers, non-repetitive propagation delays in the electronics, and ambiguities associated with time-of-arrival detection." The Ghengis robots used packetized data on spread-spectrum Proxim

[2]There are two alternate spellings for this robot: 1. "Ghenghis" as is in [5] and "Ghengis" a more common spelling found in [10]

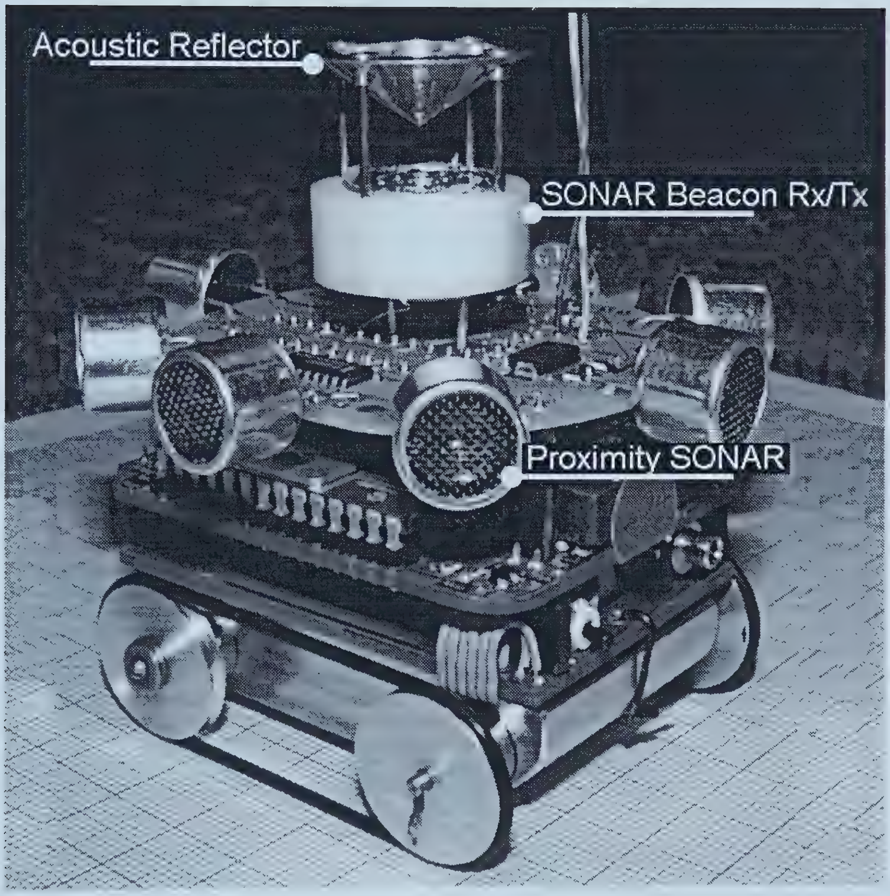

Figure 1.1: A Carnegie Mellon Millibot robot (image used with permission; labels added by the author).

Corporation radios to synchronise the SONAR beacon transmitters and the SONAR receivers on the individual robots.

Similar to the Freewave Technologies Corporation DGR-115 spread-spectrum modulated and frequency-hopping radios used by the author the radios in the Ghengis, Millibot and ROBART II [11, pages 425-26] projects are not designed for fixed latency signal transmission because of the nature of routines used to packetize and unpacketize data. In order to determine the distance between a SONAR receiver and transmitter, one must measure the time taken by the SONAR sound pulse to travel between them. To accomplish this, the transmitter and receiver circuitry must be synchronised, often by using a hard-wired connection or a radio. If a radio is used, the synchronisation signal must take a known amount of time to travel from the transmitter's processing circuitry to the receiver's processing circuitry. If this time is known the system is said to have a fixed time latency; a system which does not have a guaranteed signal propagation delay is said to have variable or unknown latency. In the case of the Freewave radios, assembling the data packets is an asynchronous process that has a theoretical minimum latency error of 10 ms. The manufacturers of other packetized radios, including the 902-928 MHz Proxim radio used by the IS Robotics researchers [5, pg. 155] and the 418 MHz Linx Technologies Corporation radios used by the Millibot team, cannot guarantee fixed latency transmission and reception. In the case of the ROBART II project, the trilateration position error due to variable radio latency was determined to be ± 21 cm.

1.3.6 SONAR Beacon Example 3: Godin at UNB

For her fourth year undergraduate University of New Brunswick (UNB) electrical engineering project [12] Alysha Godin conducted a series of interesting SONAR trilateration experiments using Polaroid 600 series transducers and 6500 series ranging boards in single echo mode. By disconnecting the trace between the transmission (XMIT) pin and the transducer on the 6500 board it became possible to disable SONAR transmission of sound by the transducer while maintaining its ability to receive signals [12, page 13]. Ms. Godin conducted experiments which determined that when the receiving transducer is rotated more than 10° from the axis of sound propagation it does not reliably detect the sound.

The research described in this thesis does employ SONAR beaconing but does not go

so far as to trilaterate position, leaving this open to future studies. Distance is measured using SONAR, but bearing (direction) is determined using a camera and related hardware and software capable of recognising a desirable target, in this case the person the robot is supposed to track.

1.4 The Research Question

This thesis addresses the initial problems facing a seismic surveying company in applying land-based semi-autonomous robotic systems to their field operations. Three specialised problems are addressed: the mechanical platform and associated electronics subsystems, a visual-only system for the potential tracking of vehicles, and a specialised SONAR and vision tracking system for use by oilfield workers. The SONAR and vision tracking system is the main topic of interest in this thesis.

1.4.1 The Pursuit of Oilfield Workers by Robots

The goal of this thesis is the demonstration of a robotic system, equipped with SONAR and vision sensors, capable of reliably following a person in a rugged outdoor environment. The research focuses on many aspects of the robot but places specific emphasis on the development and testing of the SONAR beaconing subsystem.

The introduction of a specific type of SONAR transducer which does not use acoustic reflectors as well as the use of a fixed latency (delay) radio synchronisation system are some of the interesting topics addressed.

1.4.2 Has The Question Previously Gone Unanswered?

Yes, this is the first known attempt to answer this research question. There are no known ground mobile robotic systems currently in use by the surveying industry. The author has no knowledge of any other parties engaging in the research and development of similar systems for the surveying industry, seismic or otherwise. Relevant literature has been consulted and discussions regarding the state-of-the-art have been held with people involved in the managerial, research and field operation aspects of the industry and no projects which fully address the technical issues presented in this thesis have been found.

General Robotics

No, since most, if not all, comparable systems are too expensive or too complex. Off-road robotic vehicles currently in development, including the United States Army Demo III XUV and Carnegie Mellon University Nomad projects have high mission costs. The mission costs, which place these and similar systems out of the scope of this project, are due to the following factors:

- A large and well-equipped ground crew
  - Nomad: at least a half-dozen crew
- A large Research & Development budget
  - Demo III XUV cost upwards of $24M (US) as of June, 1998
- Expensive and large number of sensors
  - Demo III Experimental Unmanned Vehicle (XUV): RADAR, passive night-vision camera, stereo camera, Forward-Looking Infrared (FLIR) camera
  - Nomad: Regular video cameras, stereo video cameras and scanning laser range-finders
- A large amount of computational power
  - Program of Intelligent Mobile Unmanned Systems (PRIMUS) robot: four PowerPC computers on-board

On-road robotic systems, including those developed by the United States Defense Advanced Research Projects Agency (DARPA) as part of the NavLab series of experiments as well as those currently being developed by automotive manufacturers are closer to the cost requirements of this particular project. Unfortunately, they all rely on environmental cues inherent to roadways.

Attempts at outdoor robotic vehicles in non-surveying fields are growing in number as time goes on. Similar to those in other industries, robotics researchers and companies have benefited from the improvement in computing power determined by Moore's Law[3]. Coupled with advances in sensor technology, more robotics projects have found themselves being applied outdoors.

[3]Moore's Law: that both transistor density of integrated circuit technology, and computer data density will double every year.

These projects have been constrained, just as any regular engineering project is. Projects requiring a high degree of autonomy such as the Nomad or Demo I, II and III XUV projects lead directly to a high mission cost because they require expensive equipment and specialised and/or numerous personnel.

Due to the cost and environmental constraints placed on the Schlumberger application semi-autonomy is clearly the only possible choice for a robotic system. Conventional tele-operation is impractical as it would require a relatively bulky interface and a dedicated operator. A system approaching full autonomy is far too expensive, complex and potentially dangerous to be fielded in non-military situations.

The environment, which is not modified for the robot prior to its arrival, has complex and dynamic constraints. This is completely the reverse of the environments in which traditional robots, like those found in factories, operate. In applications such as those addressed by Demo III a level of semi-autonomy is achieved via way-point navigation. In this particular application semi-autonomy is achieved through a passive leader, as opposed to an active tele-operator.

SONAR Localisation

Atmospheric SONAR localisation systems are not uncommon. The Ghengis 2D and ROBART II [11, pages 425-26] localisation systems offer some insight but the multiple beacons required make it unsuitable for use in this project, although the use of multiple receivers mounted on the robot for the purpose of trilaterating the position of the Pursuit Pack is a possibility for future research. As explained on page 9, the impact of this research is minimal due to the use of unknown latency, packet-based radio communication equipment.

The CMU Millibot system is also unsuitable because of the requirement for the robots' beacons to be on the same horizontal plane. This requires the robots to operate on flat surfaces, an unreasonable constraint in outdoor work. These robots also use a Linx Technologies Corporation radio synchronisation system with a non-guaranteed latency.

The UNB system offers some insight into the methods for applying Polaroid type SONAR systems, but the researcher made some underlying assumptions which cannot be made in an outdoor application. The results obtained by Godin were good but required con-

straints on the SONAR transducers which cannot be made in outdoor environments such as those encountered for this thesis. Also, Godin required that the systems be directly connected via a tether, something which is undesirable in an outdoor mobile robotic system.

This thesis describes a SONAR beaconing system which uses appropriate radio systems and a novel mixture of electrostatic and piezoelectric SONAR transducers for position determination. These are technical issues which have not been dealt with in the research projects discussed above. Between the time that the SONAR localisation system for ROBART II [11, pages 425-26] was devised in 1985 and the Millibot SONAR localisation system was developed in the late 1990's it seems that none of these projects have addressed the issue of fixed-latency radio synchronisation.

Complete Systems

Apart from the University of Alberta Polar Bear there are no readily available medium-sized rugged terrain mobile robots for conducting research relevant to this project. No commercial off-the-shelf robots provide the mobility, power or payload requirements offered by the Polar Bear. The advantages of the Polar Bear platform have been recognized by some important people involved in the development of mobile robots specifically designed for work outdoors. Mike Toscano, head of the Pentagon's Joint Project Office (JPO) was so impressed with the work accomplished by the University of Alberta team that in 1998 he extended an invitation to them to visit his staff in Washington, D.C. Benny Gothard, in charge of the Demo III XUV development at Science Applications International Corporation (SAIC) has praised the University of Alberta team for its research direction. In a speech at the (now) Intelligent Ground Vehicles Competition (IGVC) in June 1999, he emphasised to the audience how important it was for other engineering schools to follow the University of Alberta lead in outdoor mobile robot development. The Polar Bear robot won second place in the Design category at the IGVC in both 1999 and 2000. [4]

[4]The team from Virginia Tech took first place in both 1999 and 2000 by narrow margins. Bias by the American judges in this subjective category seems to be the primary factor in these decisions, especially in 1999 when the difference between Virginia Tech and the University of Alberta was approximately 1% of the maximum score. During the award ceremony in 1999 the head judge openly admitted that the decision was not based on a superior design by Virginia Tech.

1.4.3 What Makes The Question Worthwhile?

At issue is the development of a clear competitive advantage for Schlumberger Oilfield Services in the land seismic surveying business. In order to attract customers the product offered must be of high quality and low cost, and must be delivered in a timely fashion. Schlumberger currently offers a high quality surveying product. Its costs, though, are similar to its competitors because much of the work is primarily manual, labour intensive and can be performed by a relatively unspecialised work force.

By reducing time spent in the field, Schlumberger will be able to produce results faster than its competitors and will be able to reduce costs associated with field labour.

One of the problems that this thesis addresses is that the costs involved in the transportation, deployment and collection of geophysical cable are currently too high. By partially automating these tasks increased efficiency can be achieved and costs can be lowered.

From a technical point of view this thesis addresses another worthwhile issue. The author is aware of at least three ground robotics localisation projects [9] [11, page 426] [11, page 427] which use or have used separate SONAR transmitters and receivers in a beacon-like configuration for localisation. This thesis identifies and attempts to solve some of the technical challenges which those researchers did not address, including the implementation of SONAR transmitter and receiver synchronisation via radio and implementation of simplified SONAR receivers suitable for outdoor usage.

1.5 Layout of the Thesis

Chapter 1 introduces the reader to the basic research and commercial goals of this project; the current state-of-the-art is reviewed and the basic research questions are addressed. Further background information can be found in the Appendix. In Chapter 2, the reader is introduced to the Polar Bear, the outdoor mobile robot on which this research has been conducted. The two tracking systems, one solely based on vision while the other mixes vision with SONAR beaconing, are discussed in Chapter 3. Chapter 4 presents the experimental results for this project. Possible future directions for research and the conclusions reached by the author are discussed in Chapter 5.

Chapter 2

The Polar Bear Robot

2.1 Mechanical and Electrical Systems Introduction

2.1.1 The Polar Bear Robot Platform

The goal of this project has been to develop a rugged, semi-autonomous mobile robotic system. This chapter describes the basic design, construction and testing processes involved in the University of Alberta's Polar Bear mobile robot.[1]

Polar Bear's mechanical system has been designed with simplicity and ruggedness in mind, using widely available industrial components. The result is a powerful vehicle, which can be easily modified or repaired in remote locations. The four independently suspended hydraulic drive motors allow the vehicle to be skid-steered in terrain varying from ponds 50 cm. deep to asphalt and concrete.

The Polar Bear's electronics are designed to withstand rough field operation. Particular attention has been paid to the protection of circuitry from transient voltages and electromagnetic interference (EMI).

Control of the robot is achieved through an arbitration system which weighs input from the robots sensors. Direction is primarily controlled based on video data (i.e. turn left if the target is on the left side of the image, turn right if the target is on the right). Speed is controlled based on SONAR range-finding data (i.e. if an object is closer than 1 meter,

[1]The author would like to acknowledge that this section is based extensively on a series of reports submitted to the Intelligent Ground Vehicle Competition judges and on a paper submitted to the International Symposium on Robotics, 2000. Although the author was the chief contributor to these documents, Aaron Saunders, Ryan Chladny, and Jim Qualie also made strong contributions. Aaron Saunders' contribution is especially noteworthy as his primary responsibility lay in the design and construction of the mechanical system.

stop; if the way is clear in front of the robot, proceed at 90% full speed). The resulting system provides smooth control of the robot's hydraulic drive.

The Polar Bear's sturdy construction, its rugged electronics and adaptable software make it an ideal platform for the development of outdoor mobile robotics.

2.1.2 Mechanical System

The most important feature of an all-terrain vehicle platform is flexibility. Keeping the moving parts on the Polar Bear simple allowed the construction of a low maintenance, rugged robot. An air-cooled gasoline engine provides mechanical and electrical system power. This 13.4 kilowatt (18 horsepower) engine drives a hydraulic pump that pressurises the hydraulics to drive the four gear pump drive motors. This mechanical configuration allows for a powerful, simple and maintainable drive train. Through a series of trials in both the United States and Canada the Polar Bear has proven to be reliable for outdoor robotics research .

Engine

The four-stroke Robin Subaru, Inc. V2 EH65, gasoline spark ignition engine provides the mechanical power for the Polar Bear. At 44 kg. (96.9 lbs.), it is relatively lightweight and compact, measuring 0.48 m. (18.7 in.) tall, by 0.48 m. (18.8 in.) wide, by 0.32 m. (12.5 in.) deep. The V-2 configuration proved to be a very stable engine with relatively low vibration. It has a flywheel magneto that provides electrical power while the vehicle is running.

Hydraulics

Hydraulic drives reduce the complexity of power transmission and provide a low maintenance solution for industrial mobile robotics applications. Customised power train components, such as transmissions and gear trains, are generally expensive and require high maintenance. Polar Bear's hydraulic drive components are off-the-shelf and field serviceable. Parker Hannifin Corporation hydraulic components were selected due to extensive availability and use by industry. Replacing damaged hoses or wheel motors can be performed in minutes, minimising down time.

Hydraulic power is provided by a variable flow Parker pump with a built-in pressure compensator. The Robin engine drives the hydraulic pump using a flexible belt. The hydraulic flow rate can be manually adjusted at the pump, allowing the Polar Bear's maximum speed to vary from 0 to 18 kph (11 mph).

The hydraulic power to each of the four wheel motors is regulated by means of a proportional solenoid valve deck equipped with manual override levers. The valve deck's solenoids provide the interface between the mechanical and electrical systems. Bi-directional wheel motion is achieved by using simple pulse-width-modulated electronic signals. The proportional nature of the valve deck allows the implementation of traction control and low impact turning. Due to the skid-steered nature of the vehicle, a turning radius of approximately zero has been achieved.

Frame

The Polar Bear's frame is constructed of mild carbon steel. This material is both easy to fabricate and maintain, allowing for modifications and repairs to be made in the field. The bare frame is about 1.3 meters (52 in.) long, 0.4 m. (16 in.) high and 0.7 m. (28 in.) wide. With the camera mount and wheels attached, the Polar Bear is 1.7 m. (68 in.) tall. Independent suspension of each wheel is achieved by means of parallel A-arms that are connected to the frame by shoulder bolts and Beemer Precision Corporation Oilite bearings.

Wheels, Motors and Suspension

Each wheel assembly consists of two parallel A-arms connected to an aluminium Drive Motor Plate (DMP). A simple rocker arm - push rod suspension transfers the load to four compact mountain bike shocks. Wheel configuration is symmetrical about the vehicle's centre, reducing the variety of components required. A hydraulic drive motor is mounted on each DMP, providing independent four-wheel power. The resulting wheel base dimensions are 1.32 m. (52 in.) by 0.97 m. (38 in.) with a ground clear ance of 0.36 m. (14 in.).

Due to low speed requirements, complex dynamic loads are not taken into consideration; however, traversing rough terrain requires vehicle suspension to keep the wheels on the ground at all times. Inexpensive mountain bike shock absorbers with custom helical

Figure 2.1: Demonstration of the Polar Bear robot driving autonomously using a video camera to follow a white line (another possible application apart from those discussed in this thesis). (University of Alberta Experimental Farm, Edmonton, AB; 1998)

springs provide the carrying capacity for the 400 kg. (900 lb.) Polar Bear. The aluminium rocker arm provides the mechanical advantage necessary to compress the shocks, while giving each wheel eight inches of vertical travel.

Application

The design of the Polar Bear is focused on future industrial applications. The hefty steel frame, heavy hydraulic components and 18Hp gas engine were selected for compact industrial strength. The gasoline engine allows long range operation with minimal down time. Polar Bear can tow a 2700 kg. (6000 lb.) rolling weight and can carry a 140 kg. (300 lb.) payload. One third of the Polar Bear's 400 kg. (900 lbs.), is unsuspended weight located at-hub, giving it a low centre of gravity. In addition, a wide wheel base assures that Polar Bear is sure-footed on almost any terrain. The Polar Bear is capable of traversing side slopes of 25°, downhill slopes of 40°, and uphill slopes of 35°.

2.1.3 Electrical and Electronic Components

The electronic system on-board the Polar Bear is designed to be rugged, reliable and inexpensive. The components of the system include an off-the-shelf personal computer (PC), a Motorola MPC555 micro-controller, four solenoid drivers for the Polar Bear's hydraulic subsystem, six Polaroid sound navigation and ranging (SONAR) units, and one

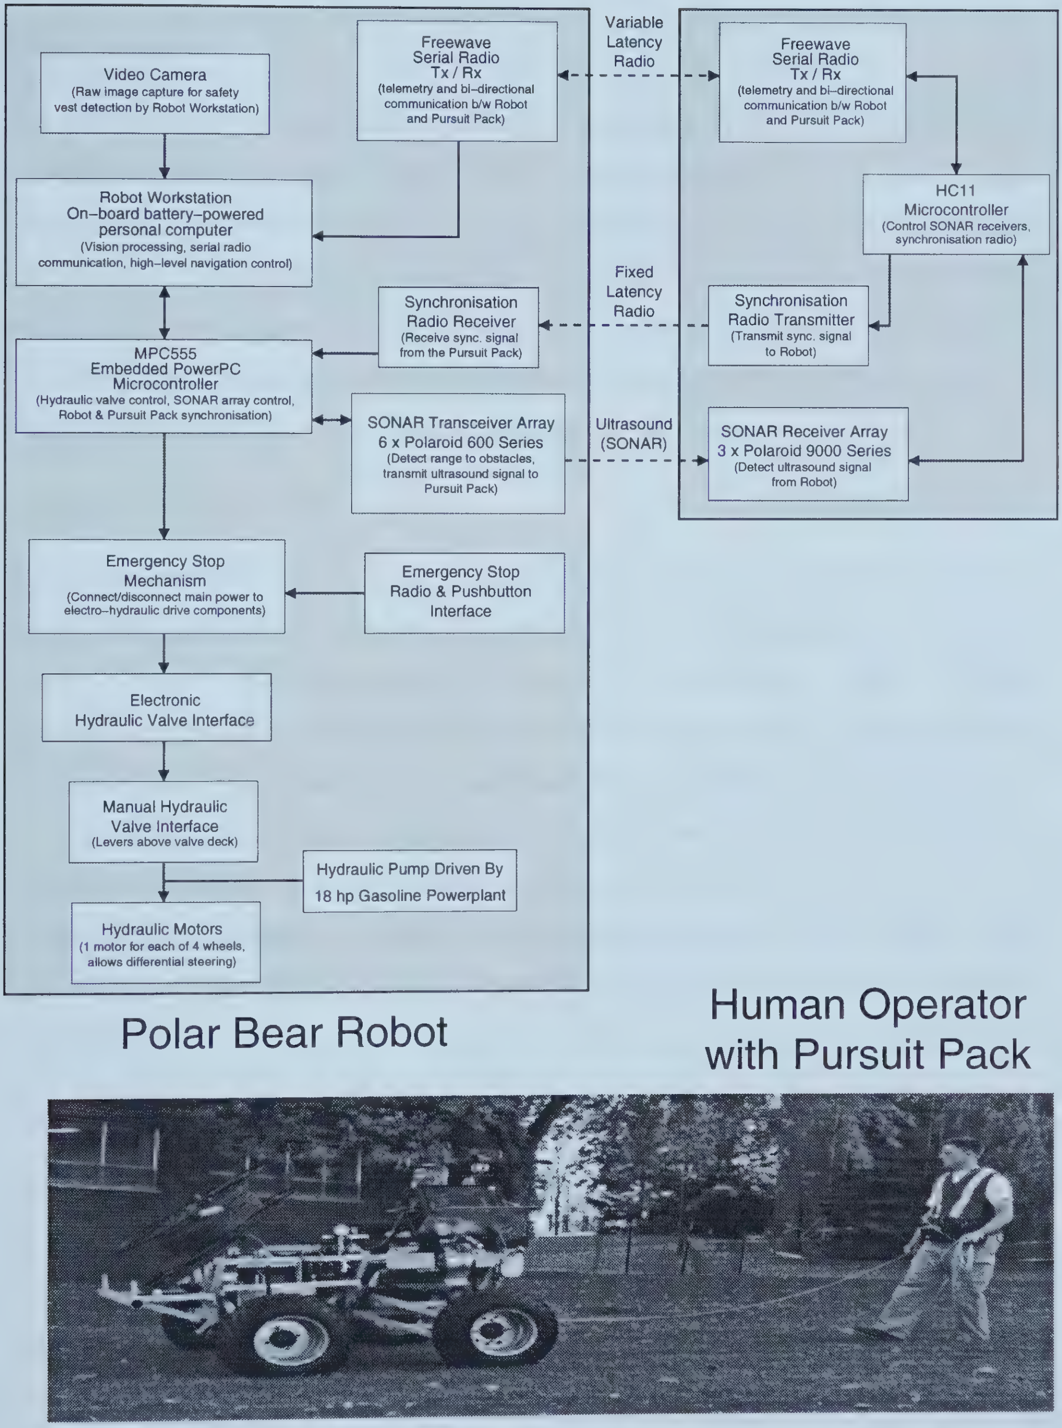

Figure 2.2: A system diagram of the Polar Bear (left) and the Pursuit Pack (right).

Sony HandyCam camcorder.

Main Computer

The main computer on board the Polar Bear is a Broadax Systems, Inc. (BSI) V8-LCD portable computer. Computer power is applied via a Keypower, Inc. 12 volt DC power supply directly connected to the Polar Bear's electrical power system. An Intel Celeron 400MHz ATX motherboard provides enough computational power for high level decision making and for video processing. Raw video data enters the computer through an off-the-shelf video capture card inside the computer which is connected to the robot's camera. The computer uses the Debian GNU/Linux operating system.

Communication and Video Capture

The main computer uses two serial ports for communication purposes. One serial port allows communication with the MPC555 low-level controller while the other is connected to the "Pursuit Pack" as is discussed in Section 3.2. Two Freewave DGR-115 spread-spectrum serial radios are used for communication between the robot's main computer and the Pursuit Pack. These radios have an advertised 32 km. (20 mile) line-of-sight range.

A modified commercial car-alarm radio transmitter and receiver pair is used on the robot and Pursuit Pack for synchronisation. A fixed 75 ms. latency between the transmitter and receiver (due primarily to signal processing) makes reliable time-of-flight measurements of robot's SONAR pulses by the Pursuit Pack SONAR receivers, as is discussed in Section 3.2.

Power and Electro-Magnetic Interference

Two lead-acid batteries and the gasoline engine's magneto power the electronics on the Polar Bear. The engine eliminates the need for a large number of bulky batteries or a recharging station. The robot's main computer can operate for up to five hours on the main battery, without requiring that the magneto be engaged. To prevent voltage dips, a recreational vehicle battery splitter ensures that when the engine is started the necessary electrical power is taken from the smaller secondary battery.

The electrical endurance of the system is estimated to be approximately 16 hours. This

is based on a current draw of approximately 20 amperes (mainly from the on-board computer), a 15 ampere charge from the engine's magneto and at least 80 amp-hours stored in the robot's lead-acid batteries. Prior to the use of a DC-DC power computer power supply, a system using a conventional personal computer AC power supply and DC/AC inverter provided an endurance of approximately 45 minutes.

If need be, the robot's batteries can be quickly recharged using an automobile and booster cables. In order to ensure the reduction of electro-magnetic interference (EMI) in the system all electrical sub-systems share a common ground point on the chassis. To further reduce EMI, shrouded resistive spark plugs and shielded power cables are used in the EH65 ignition.

Low-Level Control

The heart of the low-level control system is the MPC555 embedded PowerPC micro-controller, currently Motorola's most advanced automotive powertrain controller. The MPC555 does not require extensive shielding and is able to withstand both wide temperature ranges (-40 ° C to +125 ° C) and the EMI found near the Polar Bear's engine.

On the Polar Bear the MPC555 is responsible for executing motion control commands from the main computer, as well as control of the SONAR sensor array. These tasks are delegated to the MPC555's two Time Processing Units (TPU).

The solenoids controlling the Polar Bear's hydraulic valve deck have low current and electro-magnetic force (EMF) feedback protection requirements. Each of the Polar Bear's valve solenoid pairs draws up to 2.8 Amperes at 12 Volts DC, permitting the use of low power drivers. Filtering capacitors and transient voltage suppressors on the main power bus, and buffers between the micro-controller and the drivers ensure that voltage spikes from the Polar Bear 's engine or valve deck cannot damage the MPC555.

An array of six SONAR units, each with a maximum object detection range of almost 10 m. (30 feet), is available on the robot. These units are arranged to provide basic obstacle detection in a 120 ° arc in front of the robot.

Emergency Stop Mechanism

Two methods of stopping the robot exist which completely by-pass all other electronic subsystems. On the Polar Bear, a push-button Emergency Stop (E-stop) is located near the ignition switch. This E-stop is electrically in line with the main lead-acid battery and the hydraulic valve drive electronics. By depressing the E-Stop button an open-circuit is created and no power can reach the drive electronics, thus bringing the robot to a stop. Similarly, a radio E-Stop circuit is also placed between the main battery and the drive electronics. A remote user can bring the robot to a complete stop from up to 15 m. (50 feet) away using the radio E-stop. When the electrical power is cut off, the Polar Bear's valves return to their centre position and all hydraulic lines are closed off, immediately stopping the vehicle.

2.1.4 Image Capture

A Sony HandyCam camcorder and a PC-based image capture card are the main hardware components of the Polar Bear's vision system. Alternative video cameras, including a Hitachi security camera and a Canon infrared sensitive camera, have been tested but did not provide the image stabilisation or picture quality of the HandyCam.

Several light filters fastened in front of the HandyCam's lens provide a simple and efficient way of preprocessing the video images. A polarisation filter reduces glare and lens flares associated with reflected sunlight. Two other optical filters provide a method of reducing the direct light from the sky and enhancing contrast. These filters help to linearise the luminance and intensity of sunlight in the field of view of the camera, reducing the complexity of software processing on the images.

2.2 Control of the Hydraulic and Electric Actuators

Two types of actuators are used on the Polar Bear robot: hydraulic and electric. The four hydraulic motors, one on each wheel of the robot, are controlled by a valve deck with electronic and auxiliary manual control options. The Polar Bear's gasoline engine drives a fixed-volume pump which supplies the hydraulic fluid to the valve deck.

A Matsushita GMX-8MC045A DC brushless motor is the only electric motor on the

vehicle. It is responsible for driving the cable spooler found on the rear of the robot.

Both the hydraulic and electric motors are driven using similarly designed electronics. These designs are discussed in Section 2.3.

2.2.1 Overview of the Hydraulic Components

The spool of each valve has a pair of solenoids with spring-mounted armatures capable of regulating the direction and amount of fluid flow to a hydraulic motor. The excitation of the solenoid coil is varied by a pulse-width modulated (PWM) signal. The volume of flow is proportional to the duty cycle of the PWM signal [13].

The schematic of the hydraulic components can be seen in Figure 2.3.

When the solenoid coils are not excited the armatures are forced into a recessed position away from the valve spool by the springs. In this neutral position hydraulic fluid is circulated back into the reservoir without being directed into a motor.

When a solenoid is excited its armature pushes the valve spool, thereby opening it and allowing fluid to circulate through the hydraulic motor. The greater the excitation of the solenoid coil the further the armature travels, the larger the valve opening, and the greater is the volume of fluid which is directed to the motor.

As can be seen in Figure 2.4, the electrical interface to the hydraulic drive system is similar to the interface found on regular DC brushless motor drives. By grounding one end of each solenoid a pair of solenoids can be excited in the same manner as a DC motor.

As shown in Figure 2.6, manual control of the vehicle is possible by using the auxiliary manual levers found on the top of the valve deck.

2.3 Motor Driver Designs

Four major designs for control of the Polar Bear's hydraulic valves as well as for DC brushless motors have been implemented. The first two designs are prototypes for the second pair of designs. The second pair have a reduced number of components and are more reliable than the first pair.

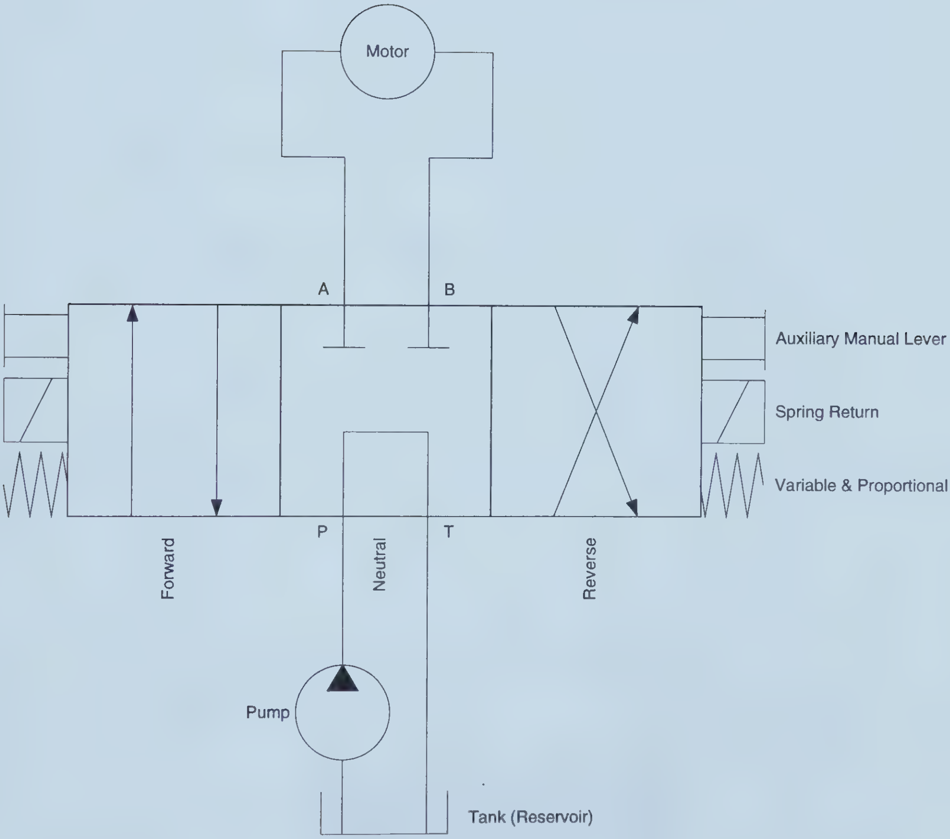

Figure 2.3: Simplified hydraulic schematic for one motor. Extra components such as the filter, radiator, and cross-over reliefs are not shown.

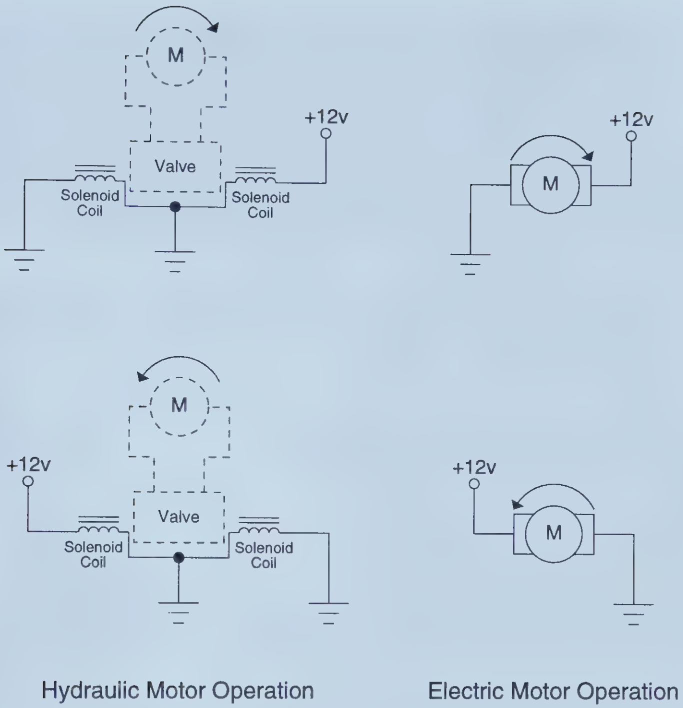

Figure 2.4: Electric operation of the hydraulic valve system versus a similar DC motor system.

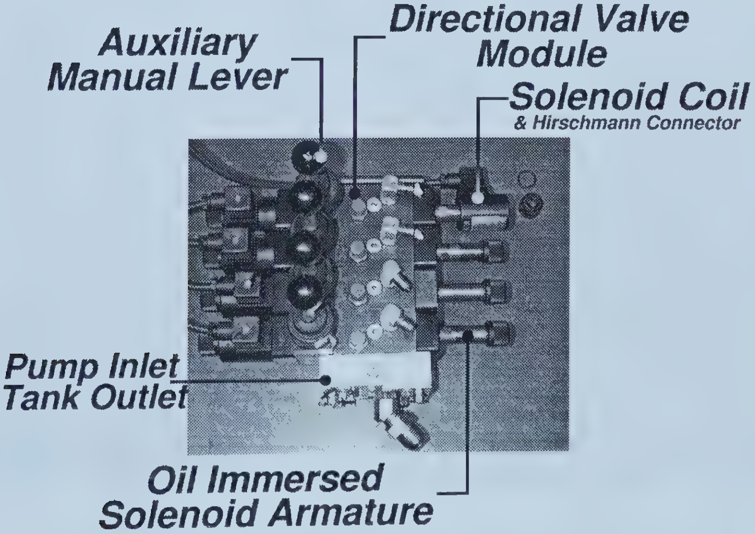

Figure 2.5: The Hydraulic Valve Deck.

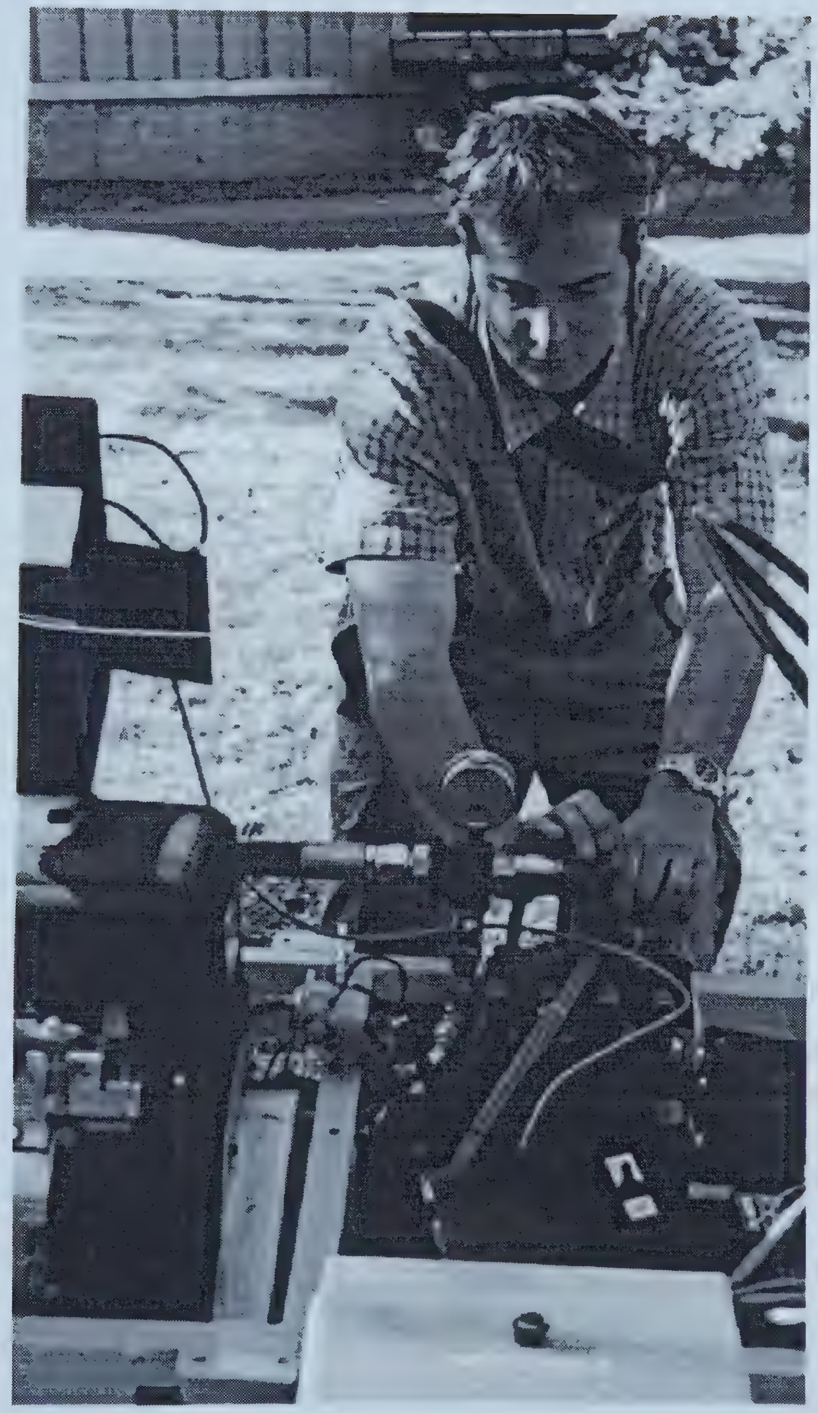

Figure 2.6: Aaron Saunders manually controlling the Polar Bear's hydraulic system (University of Alberta, Edmonton, AB; 2000).

2.3.1 Original Discrete Component H-Bridge

The first electronic interface for the hydraulic valve system was designed in early 1998. Based on a design by Loren Wyard-Scott it uses an optically isolated H-bridge configuration to drive the hydraulic valve. (The term H-bridge refers to the layout of the four power transistors in the design which resembles the letter "H".)

The n-channel Metal Oxide Semiconductor Field Effect Transistors (MOSFET) used in the design were rated for approximately 40 volts at a current draw of approximately 30 amperes. The solenoids only draw a maximum of three amperes at 12 volts. Two valve pairs, corresponding to the pair of hydraulic motors on one side of the robot, were driven by each H-bridge.

The problem of increasing the gate voltage on the two high-side MOSFETs is solved using a 9 volt battery in series with the main 12 volt battery used to drive the solenoids. With the gate-drain voltage set to 9 volts it is possible for the n-channel MOSFET to drive the high side of the solenoid load. This method of generating the necessary gate-drain voltage is used instead of the traditional charge-pump or replacement of the high-side n-channel MOSFET with an equivalent p-channel device. Although simple, this alternative does not seem to be widely used since it has not been seen in relevant literature, perhaps because of the need for regular (every few days under moderate usage) battery replacement.

2.3.2 Augmented Discrete Component H-Bridge

The objective of this design is to eliminate the 9 volt battery used in the circuit described in Section 2.3.1. The other objective is to allow a variable main power supply to be used, resulting in a board capable of driving a wider range of DC motors in addition to the solenoids on the Polar Bear. The major new component in this design is the Linear Technologies (LT) 1162 MOSFET Driver.

The LT 1162 is designed to drive n-channel MOSFETs in H-bridge applications, whether they are used in power supplies or for driving motors. The input to the LT 1162 requires a transistor-transistor logic (TTL)-compatible pulse-width-modulated (PWM) signal. External logic circuitry similar to that in the original H-Bridge design helps determine direction for the motor.

The output of the LT 1162 is divided between the low-side switch and high-side switch MOSFETs. Two charge pumps controlled by the LT 1162 through the "Boost A" and "Boost B" pins ensure that the high-side switch MOSFET gates are driven at a voltage of approximately 10 volts more than their drains.

Although initially successful, this design failed due to a large amount of voltage transients coupled between motor driver boards via the ground bus when used together; single boards did not demonstrate transients on the ground bus. In depth discussions with Linear Technologies technical support did not resolve the problem. A related MOSFET driver chip, the Harris HIP4081 was briefly experimented with in May 2000 and demonstrated some positive results but was not used due to a lack of time for rigorous testing.

2.3.3 Mixed Transistor-Relay Driver

This design, whose schematic is shown in Figure 2.7 and implementation shown in Figures 2.8 and 2.9, focuses on removing unnecessary functionality available in transistor-only H-bridge driver design. Because the objective is to control the velocity of the hydraulic motors and not their position, the rapid changes in direction possible with transistor-only design are not required.

The design uses no high-side switches, eliminating the need for expensive p-channel MOSFETs or charge-pumps for n-channel MOSFETs.

Changes in direction are achieved through a dual-pole, dual-throw (DPDT) relay. Although slower to react than a comparable H-bridge comprised of transistors, the relay is an effective solution. Since all actuators used in this project are velocity- and not position-controlled, rapid changes in direction are not required and the speed of the relay activation is sufficient. The socketed opto-isolators ensure that any electrical malfunction in the drive circuitry does not damage the micro-controller issuing the "PWM" and "Direction" signals. (An introduction to the micro-controller is found on page 21.)

Although not used in the hydraulic drive system on-board the Polar Bear this design was finally used in the spooling system. It controls the deployment of cable via a geared DC brushless motor attached to the cable spooler. It replaces a previous design which failed due to a wiring error during field trials in 1999. Ironically, the failed design was the basis for the fourth H-bridge design, a discussion of which follows in Section 2.3.4.

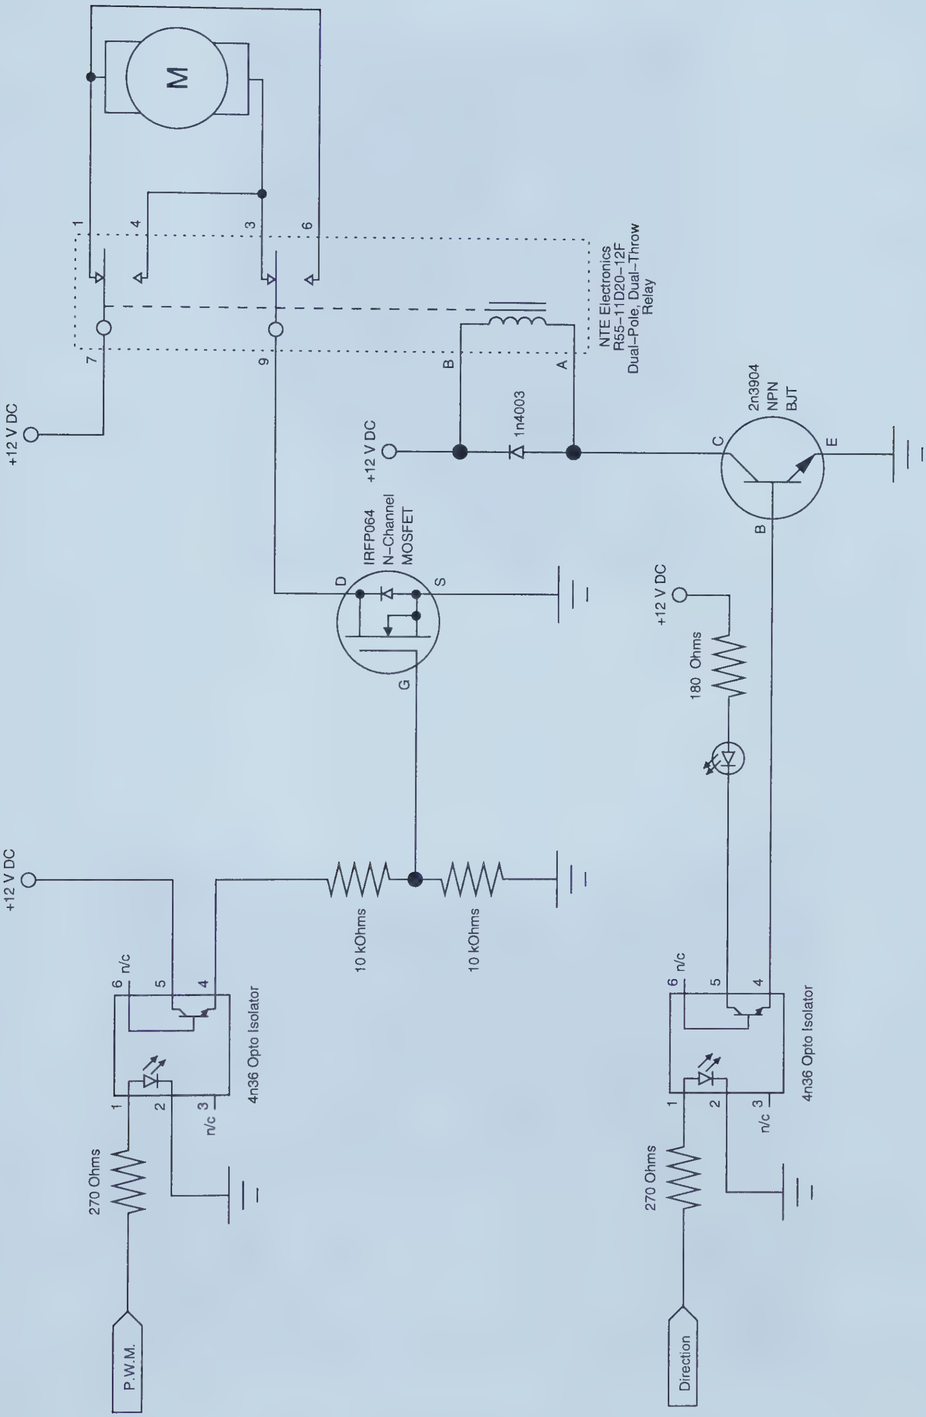

Figure 2.7: Mixed Transistor-Relay Motor Driver.

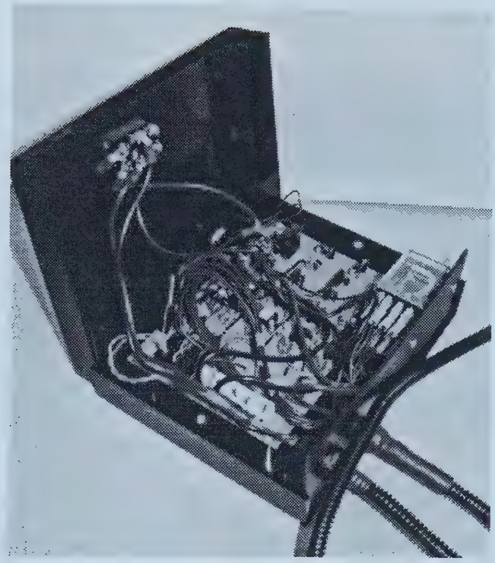

Figure 2.8: Inside Spool Control, featuring Mixed Transistor-Relay Driver.

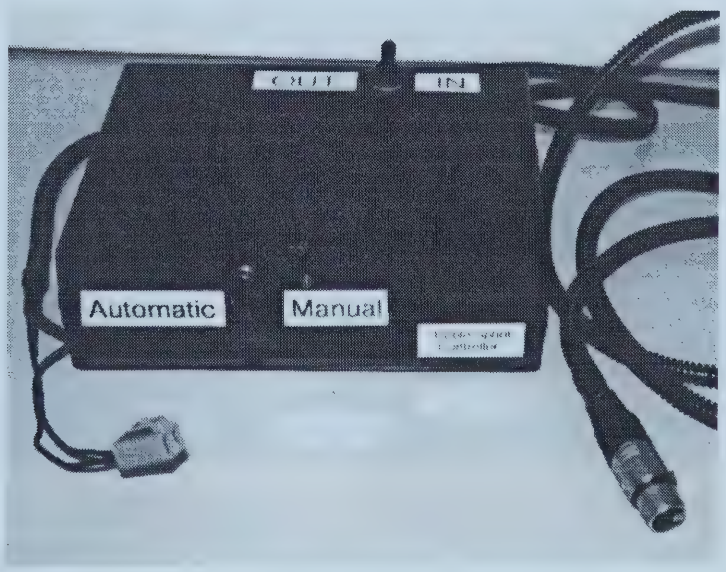

Figure 2.9: Outside Spool Control.

Figure 2.10: Aaron Saunders manually spooling geophysical cable (University of Alberta, Edmonton, AB; 2000).

Of note are the "IN OUT" and "Automatic/Manual" switches shown in Figure 2.9. These toggle switches, not shown in the Figure 2.7 for the sake of clarity, allow the Polar Bear operator to manually bypass the electronic control of the cable-spooler. This manual operation of the spool is shown in Figure 2.10.

2.3.4 Low Current H-bridge Packages

A fourth interface between the solenoid-driven hydraulics and the micro-controller was devised in June 2000. As with the other designs it relies on the paradigm that each solenoid pair can be treated, electrically, as a regular DC motor.

The driving force behind adopting this design was the need to build an inexpensive and compact circuit board. Component count is so low that it takes up half the size of the first and third designs and one-quarter the size of the second. The four STMicroelectronics L298N Dual Full Bridge Drivers cost a total of about $30 whereas the transistors used in the first or third designs are closer to $70.

The two full bridges available on each L298N are tied together so that the total maximum current output approaches four amperes, approximately two amperes more than each solenoid draws.

Unlike the designs in Sections 2.3.1, 2.3.2, and 2.3.3, this design generates a significant amount of heat. This is due to the lower current tolerance of the L298N.

A transient voltage suppressor (TransZorb) between the +12 volt and ground busses has been introduced as a safety precaution. It can dissipate large transients caused by back-EMF in functionally the same manner as a zener diode, but faster and can handle larger currents.

Both the "Direction" and "PWM" signals from the micro-controller are buffered before they reach the L298N. A bipolar junction transistor (BJT) is used to buffer the PWM signal while a second inverter available on the 7404 logic chip is used to buffer (and invert) the Direction signal.

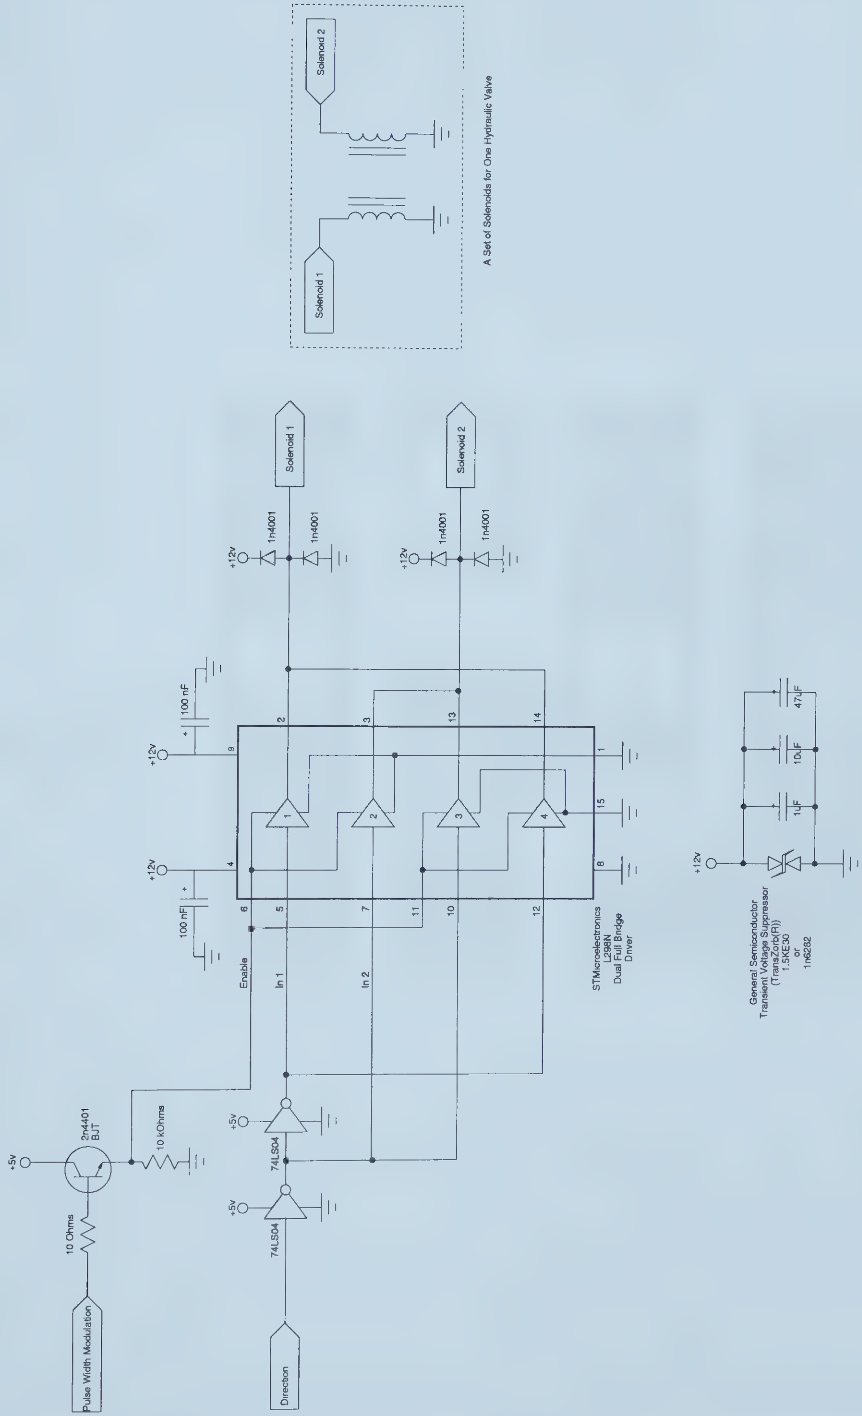

Figure 2.11: Schematic for one of four identical low current solenoid driver circuits.

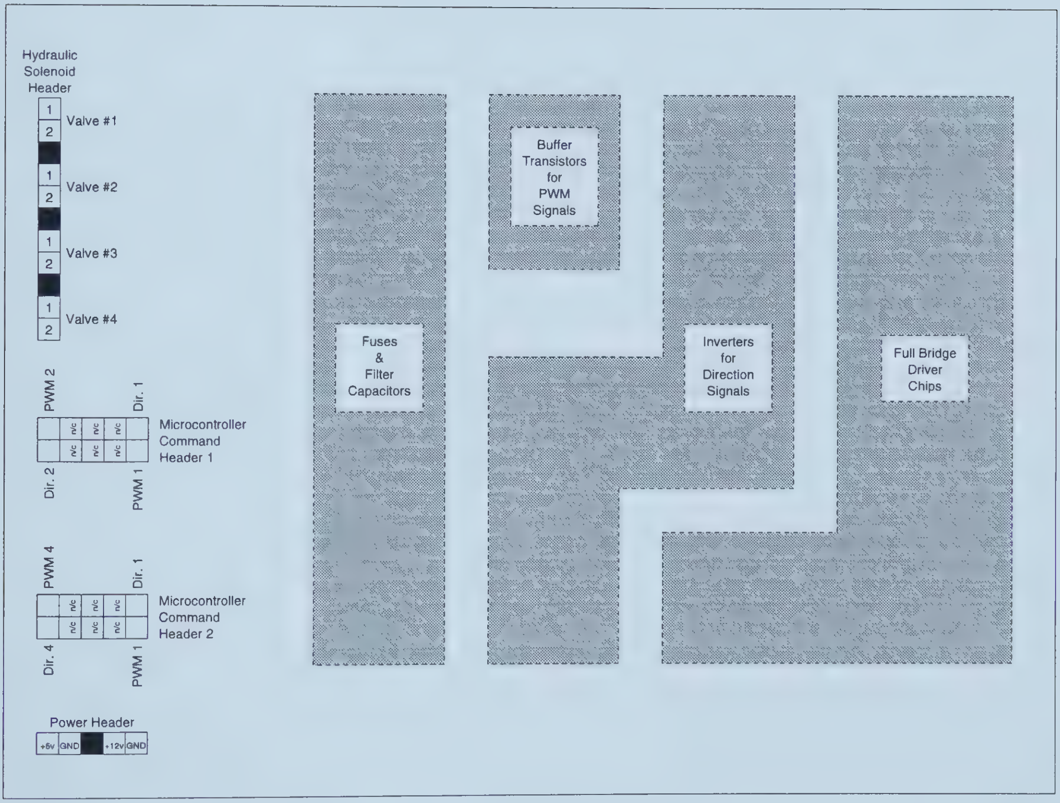

Figure 2.12: Low current driver board Layout.

2.4 Elimination of Electro-magnetic Interference

Electro-magnetic interference from the gasoline engine and the high current components on the robot was a major obstacle to the proper operation of the Polar Bear. In addition to the shielded cable and hoods placed over the engine spark plugs and leads, all electrical components on the robot share a common ground. Where possible, the grounding cable for individual components are tied to a common physical point on the robot chassis.

Chapter 3

Two Tracking Systems

In this section, a system which uses a video camera and a known target is shown to be capable of determining the three-dimensional location and pose of the target relative to the camera. This research is an alternative to the SONAR and vision tracking system discussed later on.

3.1 A Vision-only Tracking Solution

This thesis addresses two possible methods for determining range and bearing information to a target. One case involves a vision-only system capable of tracking a rigid target with known characteristics. The other involves using both vision and SONAR distance information to determine range and direction.

This vision-only system can calculate both the range and direction to the target, but also the target's *pose*. This means that it can determine the target's roll, pitch and yaw rotations in three-dimensional space.

3.1.1 Problem Definition

This section describes the preliminary study of a vision-based tracking system. The goal of the system[1] is to estimate the pose of a known object (the "target") in a laboratory environment using a stationary camera mounted on a tripod.

[1]This section is based on a project report submitted to Hong Zhang by James Andrew Smith and Andrzej Zadorozny as part of the requirements for the Computing Science 631 Introduction to Robotics course. It served as a preliminary approach to the vision localisation problem worked on by Jim Qualie for the Semi-Autonomous Guided Vehicle project.

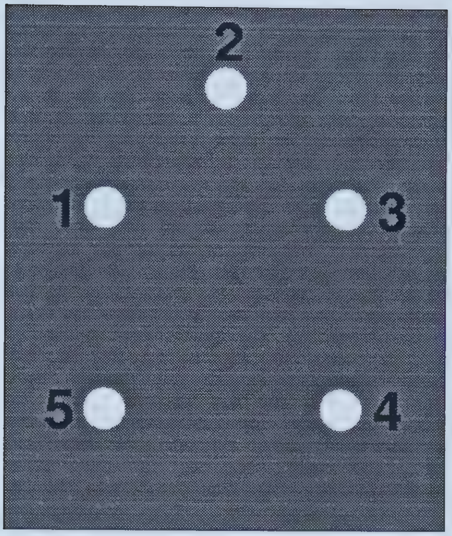

Figure 3.1: The vision-based tracking system target. The white feature points have been numbered as a convenience to the reader.

The target consists of a rigid green planar board with five white circles on it. The board measures 51 cm. wide by 76.5 cm. high. The five circles are 6.5 cm. in diameter and are positioned on the green board as shown in the Figure 3.1.

The target is designed so that the five circles, or feature points, are relatively easy for an imaging system to locate and identify. The uniform colour of the target makes it relatively easy to isolate from a cluttered image. Other arrangements of the feature points are possible. The background colour, as long as it is uniform, can also be changed.

For the trials conducted for this work the centre of the target is generally held approximately 70 cm. above and perpendicular to the ground. The target's motion is generally constrained to three degrees of freedom: two translational axes parallel to the ground (the camera's X, Z axes) and one rotational axis perpendicular to the ground (the camera's Y axis).

Only minor occlusion of the target is permitted during the experiments and attempts have been made to maintain simple lighting and background image noise conditions.

3.1.2 Solution

The objective of the project is to develop a machine vision system capable of determining the pose of a known target.

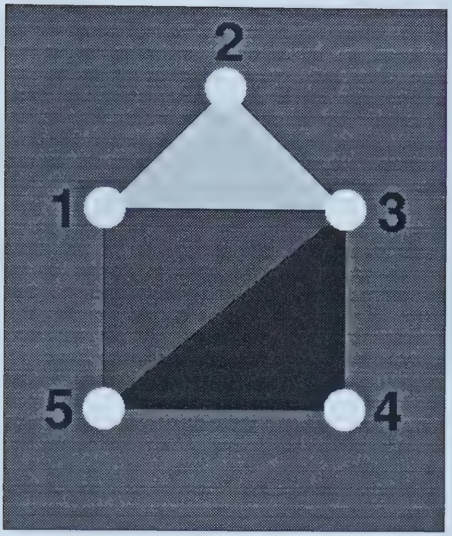

Figure 3.2: The target divided to illustrate the combination of feature points used in each of the three Perspective-3-Point iterations.

A combination of basic low-level image processing algorithms is used to remove as much unnecessary information found in the image as possible. In order to speed up the feature point extraction further image processing algorithms may be used. Cropping of the image is performed so that less time is spent on processing parts of the image that are unrelated to the target. The particular algorithms used are relatively unimportant and will not be discussed.

Once the five feature points on the target are located, three iterations of the Perspective-3-Point pose estimation algorithm are used to determine the pose of the target. Once the pose is known the locations of the target corners are calculated and then displayed on the video image. This gives a visual verification of the pose estimation in real time as the program is executing.

3.1.3 Image Processing

The video camera used in this system is an off-the-shelf Sony HandyCam. A Hauppauge WinTV capture card is used to convert the HandyCam's S-video output signal into data readable by the computer.

The raw 768 by 480 pixel image captured by the Hauppauge video capture card is converted into a 320 by 200 pixel image and is then processed before the algorithms related

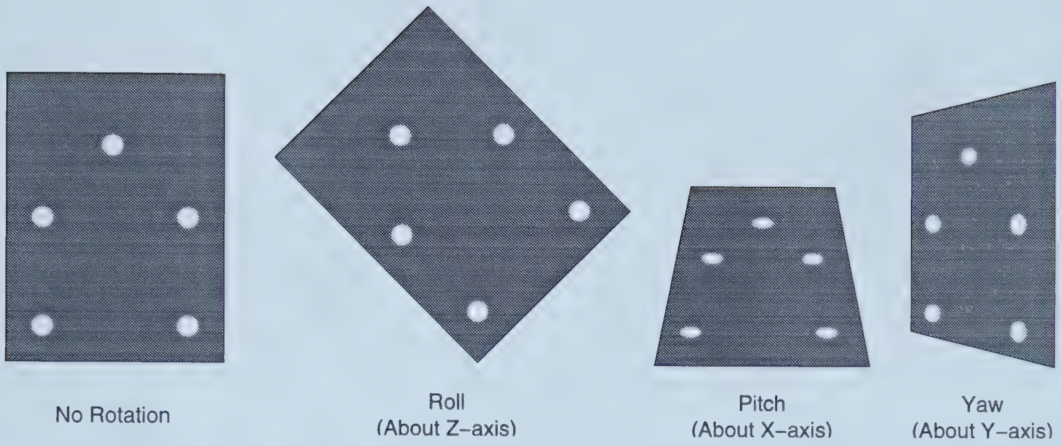

Figure 3.3: Illustrations of the three possible single rotations of the target.

to pose estimation are performed. The colour information is stored and used throughout the image processing as it contains valuable information about the scene and the target. Using this colour information it becomes possible to distinguish the target from the background, as well as the feature points on the target. The white feature points are isolated from the green background colour of the target by an algorithm which scans horizontally across the image, looking for blocks of white pixels surrounded by green ones.

3.1.4 Pose Estimation

The estimation of the pose of an object is the determination of the translation of a given point on a target as well as the rotation of the target in relation to a known frame of reference. In the case of this project the known frame of reference is the camera's coordinate frame and the target is a green board with white circular feature points.

Perspective-3-Point Algorithm

The Perspective-3-Point (P-3-P) algorithm, described in its general Perspective-n-Point (P-n-P) form in [4] and [2], uses the Law of Cosines to determine the straight line distance from the origin of the camera's frame of reference to a feature point on the target.

To be successful, the algorithm must be told the real-world distance between each of the feature points on the target. Other required information include the camera focal length, the number of horizontal and vertical pixels, the vertical and horizontal sizes of the pixels

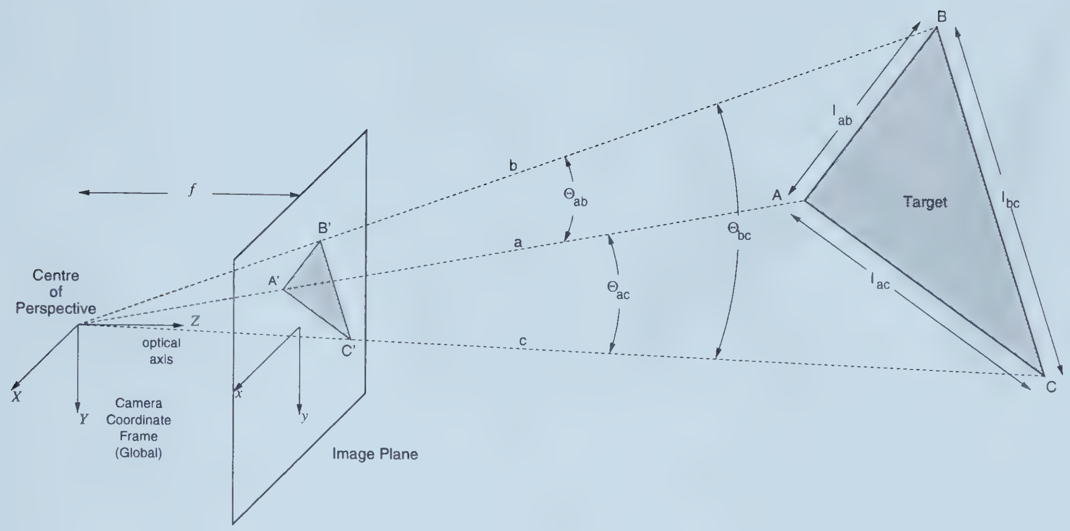

Figure 3.4: Variables used in Perspective-3-Point calculations.

and the pixel coordinates of the three feature points.

Before the Perspective-3-Point algorithm is implemented, some manipulation of the coordinate system is required. First, the pixel coordinate system is changed to image coordinates. Then, the frame of reference for the feature points on the charge-coupled display (CCD) image plane is converted so that it is no longer the upper left corner of the image[2], but rather is the centre of the CCD. As well, the angle required by the Cosine Law is calculated using the distance between the feature points on the image plane as well as the focal length of the camera.[3]

The Perspective-3-Point algorithm code sets up the five coefficients for a quartic[4] equation. The coefficients are derived from [2] and the relevant errata [3].

The quartic equation is solved using a subroutine from [14]. Multiple solutions for each quartic equation are returned from this subroutine.

It is commonly known that multiple solutions for the distance to the feature points result if only three feature points are used. From [15, notes dated March 10], where n equals the number of feature points on the target,

[2]This is from the perspective of a virtual image plane, consistent with how the user sees the image in the video camera view-finder or on a computer monitor. The focal length is assumed to be positive.

[3]The basic calculations are based on the pin-hole camera model.

[4]Quartic: a fourth order polynomial.

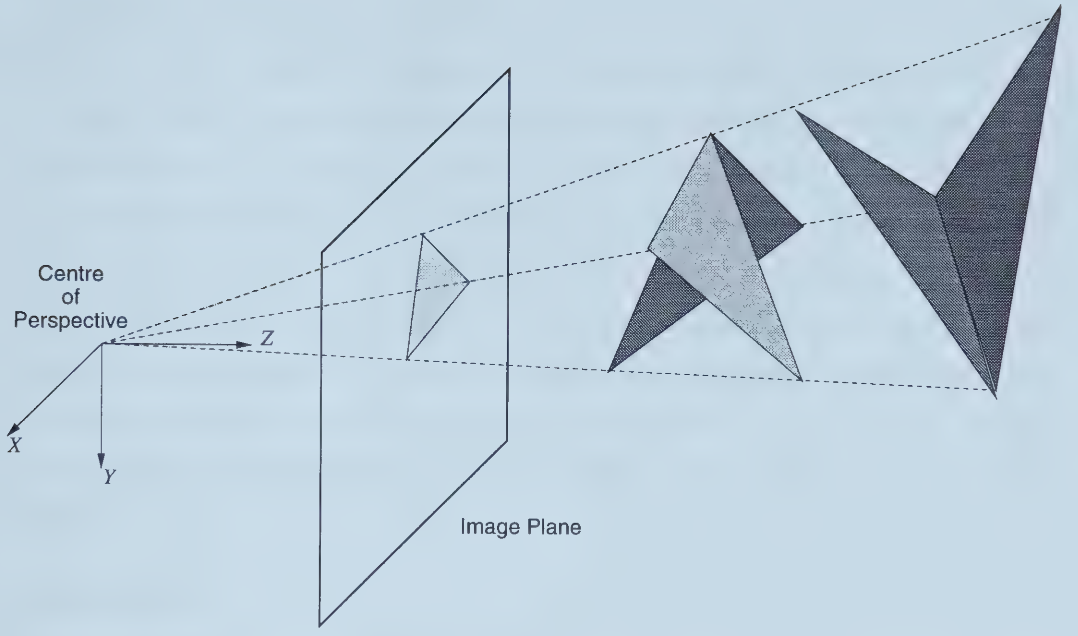

Figure 3.5: A target with three feature points (n=3) has four physically realisable pose solutions. In other words, four different triangle poses will generate the same image in the camera.

- If $n < 3$ there are infinite pose estimation solutions.
- If $n = 3$ there are eight pose estimation solutions, of which only four are physically possible.
- If $n = 4$ there are two pose estimation solutions if the feature points are not coplanar; one solution is possible if the points are coplanar.
- If $n > 4$ there is always a unique pose estimation solution.

By applying the Perspective-3-Point solution three times to the target, using three different combinations of the five feature points, as in Figure 3.1, it becomes possible to ensure a unique solution. The advantage of this method is that it uses multiple iterations of the proven P-3-P algorithm whereas the P-4-P or P-5-P equivalents, although probably computationally more efficient, would need to be redeveloped since [4] is riddled with errors.

By comparing the results from different P-3-P calculations which use common feature points it becomes possible to determine a unique solution.

Specifically, in the case of the target used in this project shown in Figure 3.2, the Perspective-3-Point algorithm is applied to three triangles formed by feature points (1,2,3), (1,3,5) and (3,4,5). Using solutions for the quartic equations describing the distances to the points in triangles (1,2,3) and (3,4,5) are used to determine the correct solution for (1,3,5).

In practice, the method described previously must be modified slightly. Because the calculated distances from the camera's centre of perspective to feature points three and five are not exactly the same (due to factors such as misalignment of target feature points, variations in feature point size or shape, and camera lens distortion) for triangle (1,3,5) as they are for triangle (3,4,5) the algorithm must look for results with the smallest variance between the distances found for the common feature points. Using this method a unique solution is found.

Pose Estimation

Pose estimation is accomplished via a few easy steps. Once the three dimensional coordinates of feature points one, three and five are known the three dimensional space coordinates of the "centroid" of the target are calculated. (In fact, this is not the real centroid but rather a centroid of convenience.) This centroid is calculated to be the point midway between feature points three and five. Knowing the location of the centroid essentially yields the translation of the target from the origin of the camera's reference frame.

To solve for the rotation of the target about the camera's frame of reference the three-dimensional space coordinates of each of the feature points are translated to the origin by the distance calculated for the centroid. Once translated the X (pitch), Y (yaw) and Z (roll) component angles are calculated between the centroid and feature point three. Using an arctangent function (*atan2()*) and the component distances between the origin and feature point three it becomes possible to determine the rotations about the X (pitch), Y (yaw) and Z (roll) axes.

$$Roll = atan2(Y_{fp3} - Y_{centroid}, Z_{fp3} - Z_{centroid}) \tag{3.1}$$

$$Pitch = atan2(X_{fp3} - X_{centroid}, Z_{fp3} - Z_{centroid}) \tag{3.2}$$

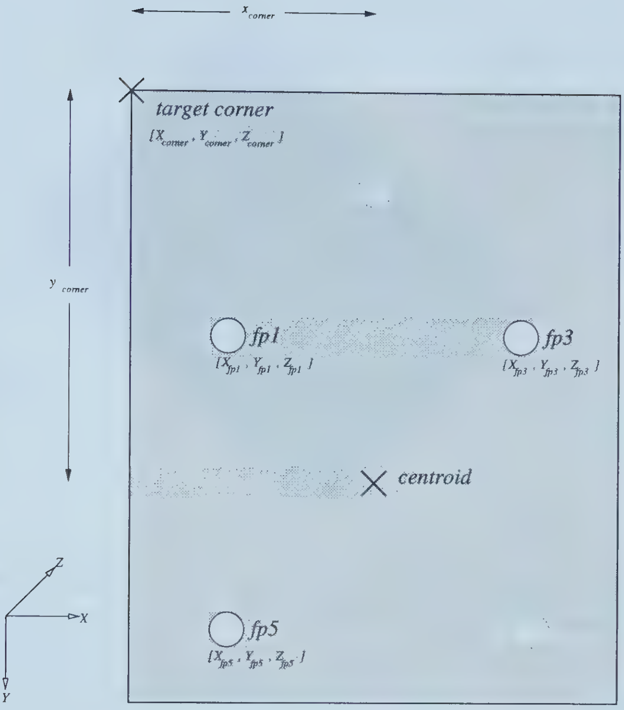

Figure 3.6: The important variables on the target which are used in pose verification.

$$Yaw = atan2(Y_{fp3} - Y_{centroid}, X_{fp3} - X_{centroid}) \tag{3.3}$$

Pose Verification

To demonstrate that the correct pose has been calculated, the three dimensional space locations of the four corners of the green board are estimated. Once the three-dimensional location of a corner is known its image plane coordinates are calculated. Lines are drawn to outline the estimated edge of the target and the result is superimposed on the actual video image, as seen in Figure 3.9.

Please refer to Figure 3.6 for a visual interpretation some of what is described below. Figure 3.4 illustrates how a generic target relates to the video camera and should also be referred to.

The global coordinates (whose origin is on the CCD of the video camera) of three of the five feature points are used for pose verification. In addition, the location of the target corners relative to the centroid of the target are used. The centroid, the midway point

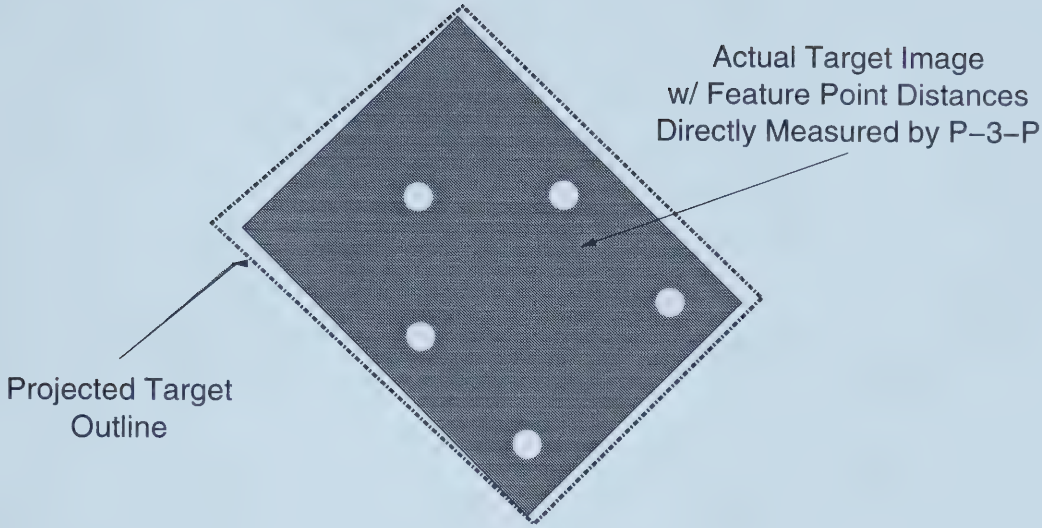

Figure 3.7: Illustration of pose verification. Once the distance to the feature points is known via P-3-P, the resulting pose is used to estimate the location of the target's four corners. The closer the box (dashed line in this figure) joining the estimated corner locations is to the outline of the target, the better the estimated pose is.

between feature points three and five, is considered to be the origin of the target's local coordinate system. The locations of the target's corners are initially defined in terms of this local coordinate system. The result of the pose verification will yield their estimated locations in the global coordinate system.

Global coordinate variables are defined in upper-case (e.g. Y_{corner}). Local coordinate variables are defined in lower-case (e.g. y_{corner}).

$$X_{horizontal} = \frac{X_{fp3} - X_{fp1}}{\sqrt{(X_{fp3} - X_{fp1})^2 + (Y_{fp3} - Y_{fp1})^2 + (Z_{fp3} - Z_{fp1})^2}} \tag{3.4}$$

$$X_{vertical} = \frac{X_{fp5} - X_{fp1}}{\sqrt{(X_{fp5} - X_{fp1})^2 + (Y_{fp5} - Y_{fp1})^2 + (Z_{fp5} - Z_{fp1})^2}} \tag{3.5}$$

Here, the global coordinates for feature points one, three and five are used. These points are used since the line joining feature points one and three is parallel to x_{corner} while that joining one and five is parallel to y_{corner}. The line joining feature points one and three is considered to be horizontal and that between one and five, vertical, hence the subscripts on the left-hand side of Equations 3.5 and 3.4.

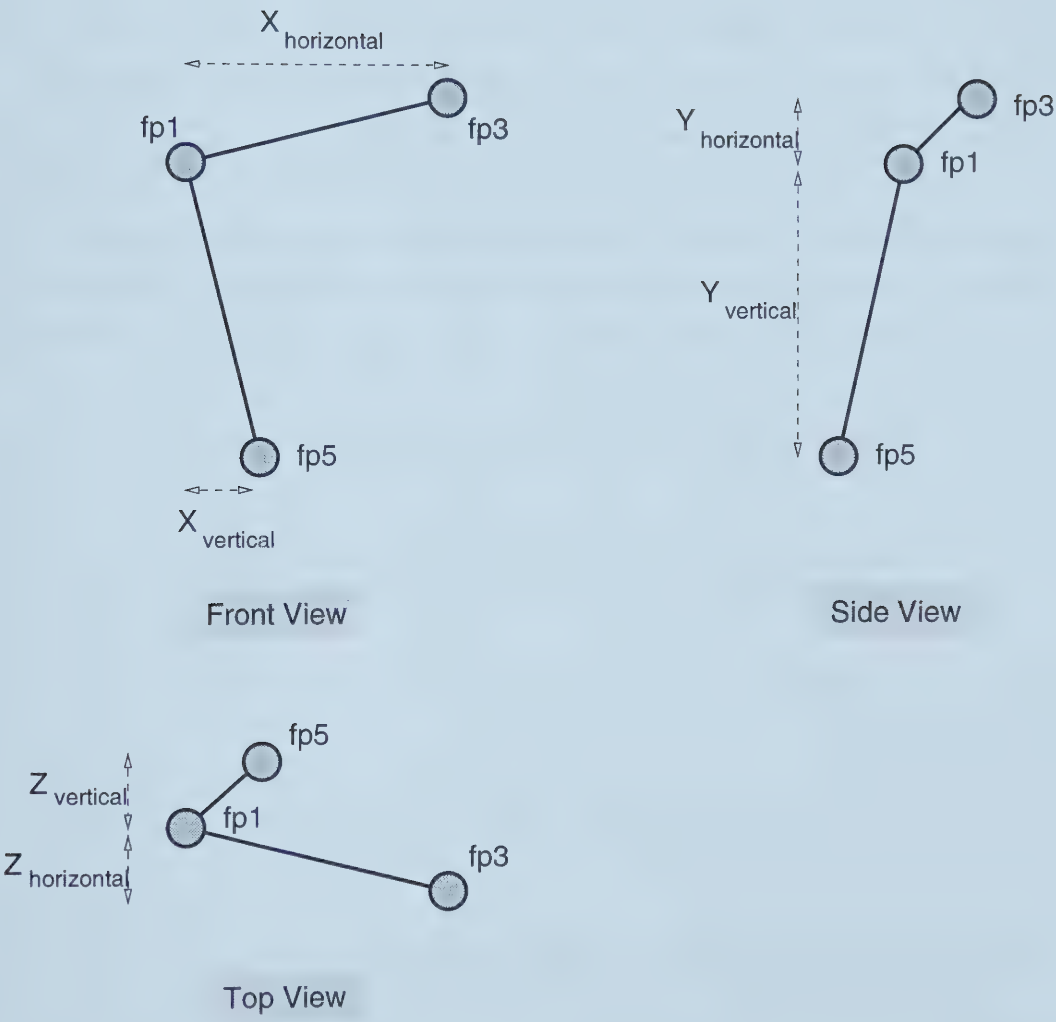

Figure 3.8: Here, some of the variables used in pose estimation are shown. Three feature points (out of a possible five) are shown from three different perspectives, front, top and side.

Note that the denominators in both equations yield the length of the line joining two feature points. Equation 3.5 yields the sine of the angle about the negative Z-axis that the target is rotated. Similarly, Equation 3.4 yields the cosine of that angle.

$$X_{corner} = X_{centroid} + x_{corner} * X_{horizontal} + y_{corner} * X_{vertical} \tag{3.6}$$

Equation 3.6 combines the local coordinate distance between the centroid and the corner (x_{corner} and y_{corner}) with the results of equations 3.5 and 3.4. The result is added to the global coordinate X value for the target's centroid.

A similar set of equations, described in equations 3.8 through 3.12, yield similar results for the Y and Z components of the target corner. The equations 3.8 and 3.7 apply to rotations about the X-axis, while 3.11 and 3.10 refer to those about the Y-axis.

$$Y_{horizontal} = \frac{Y_{fp3} - Y_{fp1}}{\sqrt{(X_{fp3} - X_{fp1})^2 + (Y_{fp3} - Y_{fp1})^2 + (Z_{fp3} - Z_{fp1})^2}} \tag{3.7}$$

$$Y_{vertical} = \frac{Y_{fp5} - Y_{fp1}}{\sqrt{(X_{fp5} - X_{fp1})^2 + (Y_{fp5} - Y_{fp1})^2 + (Z_{fp5} - Z_{fp1})^2}} \tag{3.8}$$

$$Y_{corner} = Y_{centroid} + x_{corner} * Y_{horizontal} + y_{corner} * Y_{vertical} \tag{3.9}$$

$$Z_{horizontal} = \frac{Z_{fp3} - Z_{fp1}}{\sqrt{(X_{fp3} - X_{fp1})^2 + (Y_{fp3} - Y_{fp1})^2 + (Z_{fp3} - Z_{fp1})^2}} \tag{3.10}$$

$$Z_{vertical} = \frac{Z_{fp5} - Z_{fp1}}{\sqrt{(X_{fp5} - X_{fp1})^2 + (Y_{fp5} - Y_{fp1})^2 + (Z_{fp5} - Z_{fp1})^2}} \tag{3.11}$$

$$Z_{corner} = Z_{centroid} + x_{corner} * Z_{horizontal} + y_{corner} * Z_{vertical} \tag{3.12}$$

From these values for the three-dimensional location of the corner, the corner's two-dimensional pixel coordinates $[u, v]$ on the image plane needs to be determined.

$$u = (\frac{f * X_{corner}}{k_u * Z_{corner}}) * u_c \tag{3.13}$$

$$v = (\frac{f * Y_{corner}}{k_v * Z_{corner}}) * v_c \tag{3.14}$$

In these equations, f refers to the camera's focal length; X_{corner}, Y_{corner}, Z_{corner} refer to the three-dimensional space X, Y, Z coordinates, respectively; k_u and k_v refer to the horizontal and vertical CCD pixel dimensions; u_c and v_c refer to the horizontal and vertical coordinates of the location of the CCD's centre pixel.

Algorithm & Camera Parameter Verification

Placing the camera on a tripod and the target on a wall, surveying string and levels were used to determine the straight-line distance between the camera's focal point and the three feature points of the known target. Using the camera parameters from [16] and assuming that the focal length of the camera was the specified 0.0041 meters the P-3-P algorithm was run using the measured distances. This resulted in two possible configurations, one of which estimated the feature point distances to within 7% of the actual measured values, the other was not representative of the actual target.

The experiment was repeated and the results were consistently off by between 5% and 8%. The focal length parameter was adjusted until the P-3-P trials yielded less errors. The final focal length (with the camera set to its widest field-of-view) was experimentally determined to be 0.0049 m.

3.1.5 The User Interface

The system outputs results both graphically and textually, as can be seen in Figure 3.9. The textual information includes the size of the image, the current frame number, the red-green-blue (RGB) values of the target, the three-dimensional locations of the target feature points and corners, and the target's translation and rotation (i.e. the target's pose). All coordinates are with respect to the camera's frame of reference.

Graphical information is displayed in a 640 by 400 pixel window, with a status bar below it. The five feature points and edge of the target are highlighted, demonstrating the calculated target pose. The five feature points identified by the cross-hairs are a result of the automatic feature detection performed during the image processing. The target's highlighted edge is a result of the pose estimation and three-dimensional perspective projection.

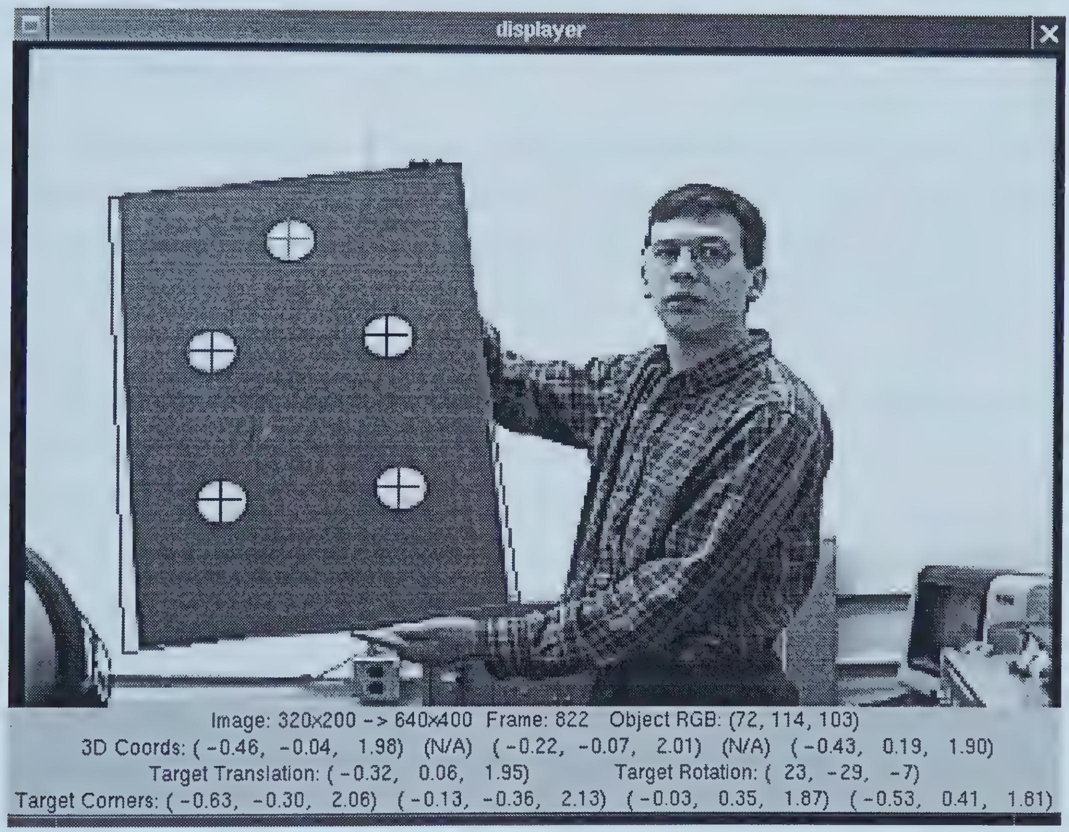

Figure 3.9: Screen capture of target tracking software. Note the addition of the feature point cross-hairs and the lines along the outside of the target demonstrating estimated pose.

On a Intel Pentium II 300 MHz computer the image capture and processing system occurs at 4.25 frames per second. This could be increased to approximately 8 frames per second if some superfluous processes (e.g. the ones related to ethernet image transmission) are removed. This frame rate was demonstrated by Jim Qualie in a related project (the visual detection of the safety vest).

3.2 A SONAR & Vision Tracker with Pursuit Capability

3.2.1 Introduction

The main goal of this research is the initial development of a complete mobile robotic system capable of following a person outdoors. Once this task is properly implemented the research should shift focus to the application of the robot to specific outdoor tasks; that is, to bring the project to Levels 7 through 9 on the NASA Technology Readiness Levels scale.

The problem of the robot (in this case, the University of Alberta Polar Bear) following a person has been broken up into two separate sub-problems: obtaining the person's distance and obtaining the person's bearing relative to the front of the robot. The robot's navigational goals are modified based on these two data. It speeds up when the person is too far away (currently 3 m.), and slows down and stops when the person is too close (currently 1 m.). The robot also modifies its direction based on where the person is in the robot's visual field.

Since most Schlumberger employees are required to wear bright safety vests while working in the field it was decided that the robot should take advantage of this fact. Also, it was decided that any additional pieces of equipment that the person would have to wear would have to be small, light and relatively comfortable to wear.

3.2.2 Determination of Bearing using Vision

Industrial safety vests come in a wide variety of colours, shapes and sizes. The shape of any given vest changes based on what is worn underneath, as well as on the body type of the wearer. These conditions make *a priori*[5] computer models of the vest impractical, if not impossible, for the robot to use. Instead, it was decided that the vision system should

[5]a priori: Latin for "to know ahead of time."

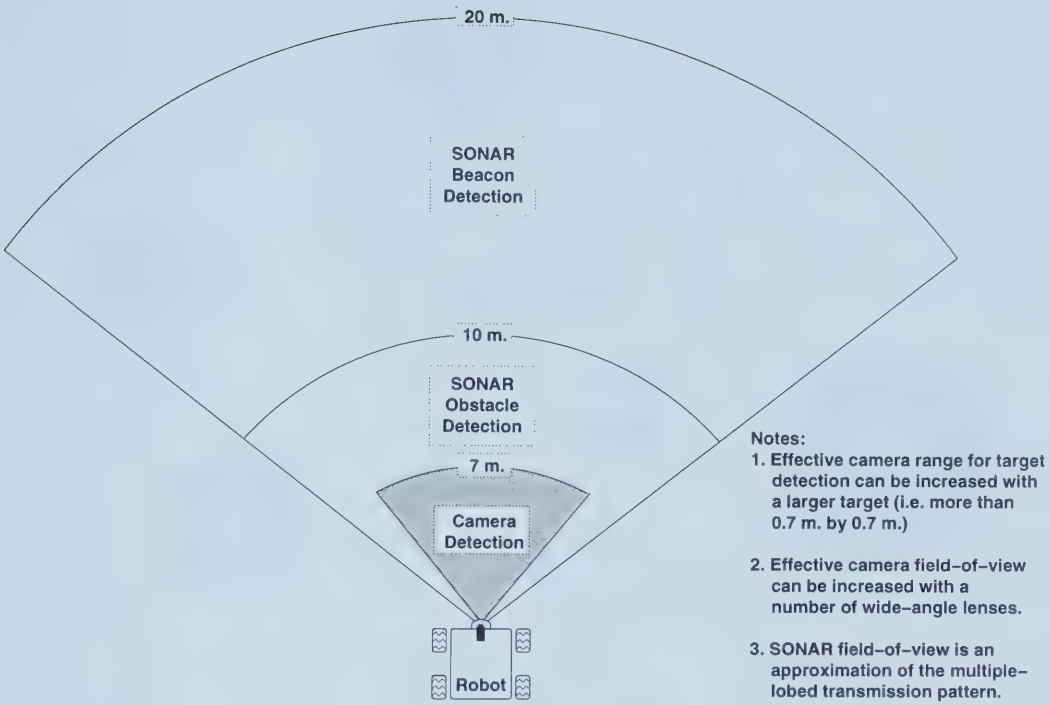

Figure 3.10: The range and fields-of-view of the video camera and SONAR sensors.

only rely on the fact that the vest is of a relatively uniform, distinct and unknown colour. The colour is to be selected by a human operator using the personal computer mounted on the robot while in the field.

The vision system, an off-the-shelf video camera interfaced to the on-board personal computer, returns the bearing of the person relative to the camera by isolating the colour of the vest in the image. Bearing is estimated based on the size and position of the isolated vest. This system, developed by Jim Qualie[6], uses a base set of code developed Jason Gunthorpe's[7] and other open source programmers. This same base set of code provided the fundamental camera-computer software interface for the work described in Section 3.1.

3.2.3 Range-Finding using a SONAR Array

Several methods exist for determining range in robotics. SICK Corporation sells a relatively expensive LASER range-finder [17] which is commonly used by industry and the

[6]Jim Qualie has been involved in both the ARVP and SAGV, primarily with respect to the development of vision capability

[7]Jason Gunthorpe is responsible for much of the programming work for the ARVP

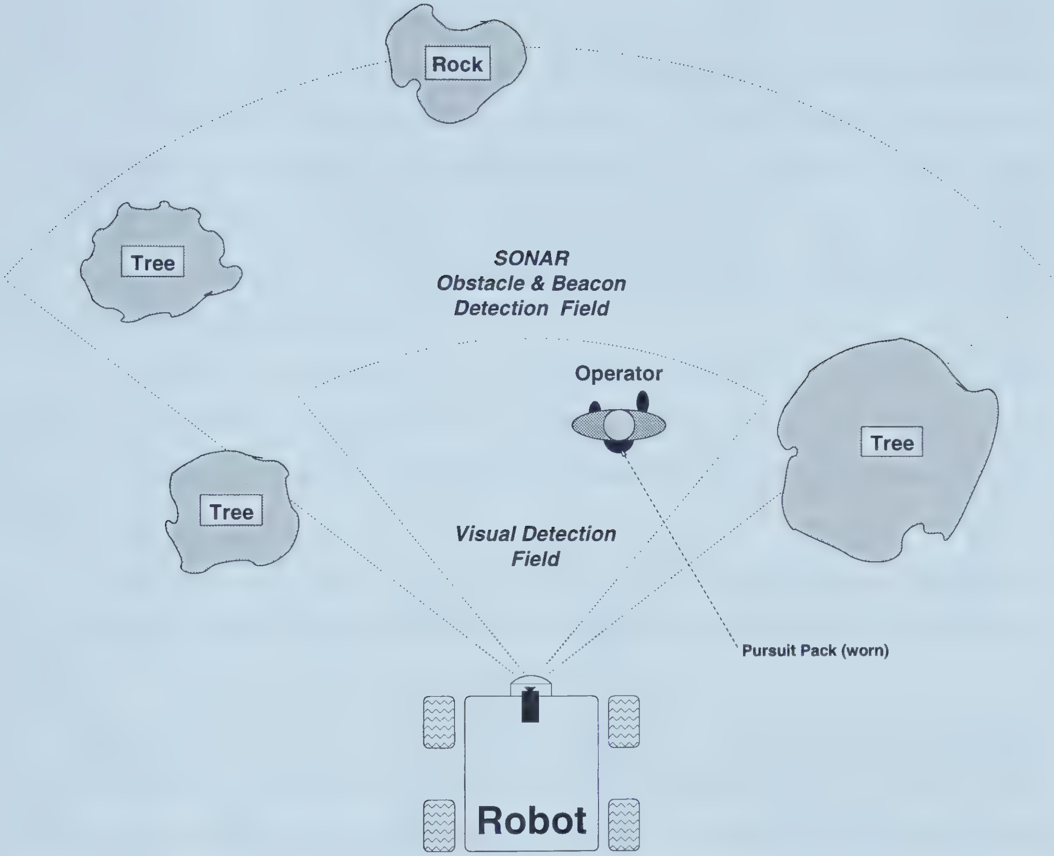

Figure 3.11: Bird's eye view of operator leading the robot between obstacles.

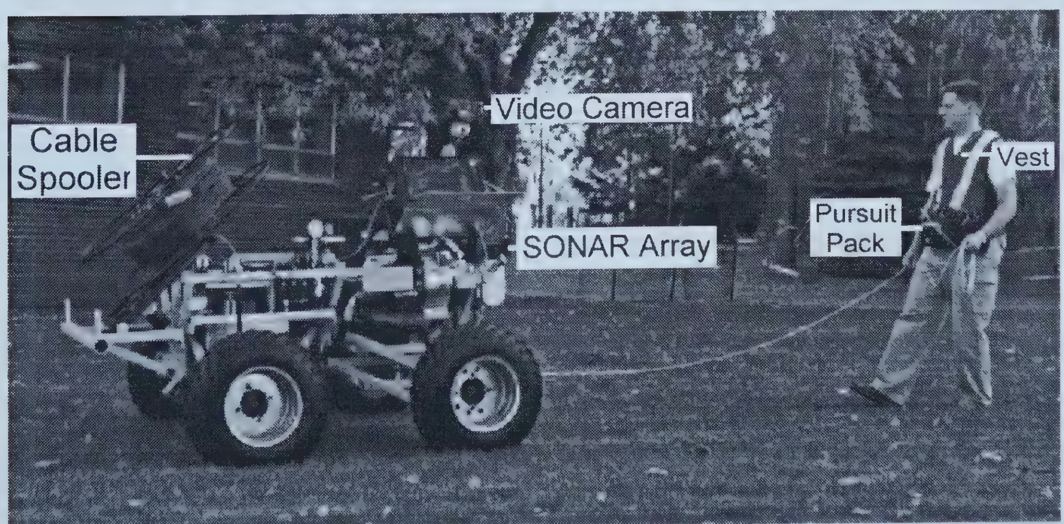

Figure 3.12: Polar Bear pursuing Jim Qualie. Note the safety vest and the Pursuit Pack worn about his waist.(University of Alberta, Edmonton, AB; 2000)

research community. Modulated infrared systems sold by robotics hobbyist distributors such as Calgary's HVW Technologies, Inc. are a cheaper but shorter range alternative. Stereo video systems, sold by a number of companies, including Calgary's VX Optronics Corporation and Vancouver's Point Grey Research, Inc., or built from discrete camera components and manually calibrated are also available. SONAR, using widely available kits and components from Polaroid Corporation, Senix Corporation or Australia's Hexamite.com, is another possible solution.

Since a model of the dimensions of the safety vest cannot be known and a suitable stereoscopic camera is not available to the researchers another method for determining the range to the person is required. SONAR is often used in robotics applications for range-finding tasks and has been selected for this one, as well.

Traditionally, mobile robots use simple SONAR transducers to range-find objects in their immediate surroundings. Unfortunately this generally provides little information regarding the nature of the object other than its relative distance.

A small electronics and SONAR transducer package, named the Pursuit Pack, has been developed to aid in the location of the person. It is worn about the waist of the person to be pursued. It is used to measure the time-of-flight of the ultrasonic SONAR pulses generated by the robot's SONAR array when they encounter the person. These measurements are converted into distances and transmitted back to the robot via a radio link, thus informing it of the distance between it and the Pursuit Pack. Further details about the Pursuit Pack may be found in Section 3.2.5. The reader should refer to Figures 3.16, 3.21 and 3.22 for visual representations of the Pursuit Pack.

Other SONAR transmitter and receiver configurations are possible. One possibility, similar to trilateration method used in both [9] and [12] involves a transmitter on the person and multiple SONAR receivers on the robot. The advantage to such a configuration is that the trilaterated position of the person could be determined using three or more receivers on the robot. Unfortunately, since all the SONAR transducers currently used in this system operate at about 50 kHz crosstalk between robot transmitters and Pursuit Pack transmitters would become a concern.

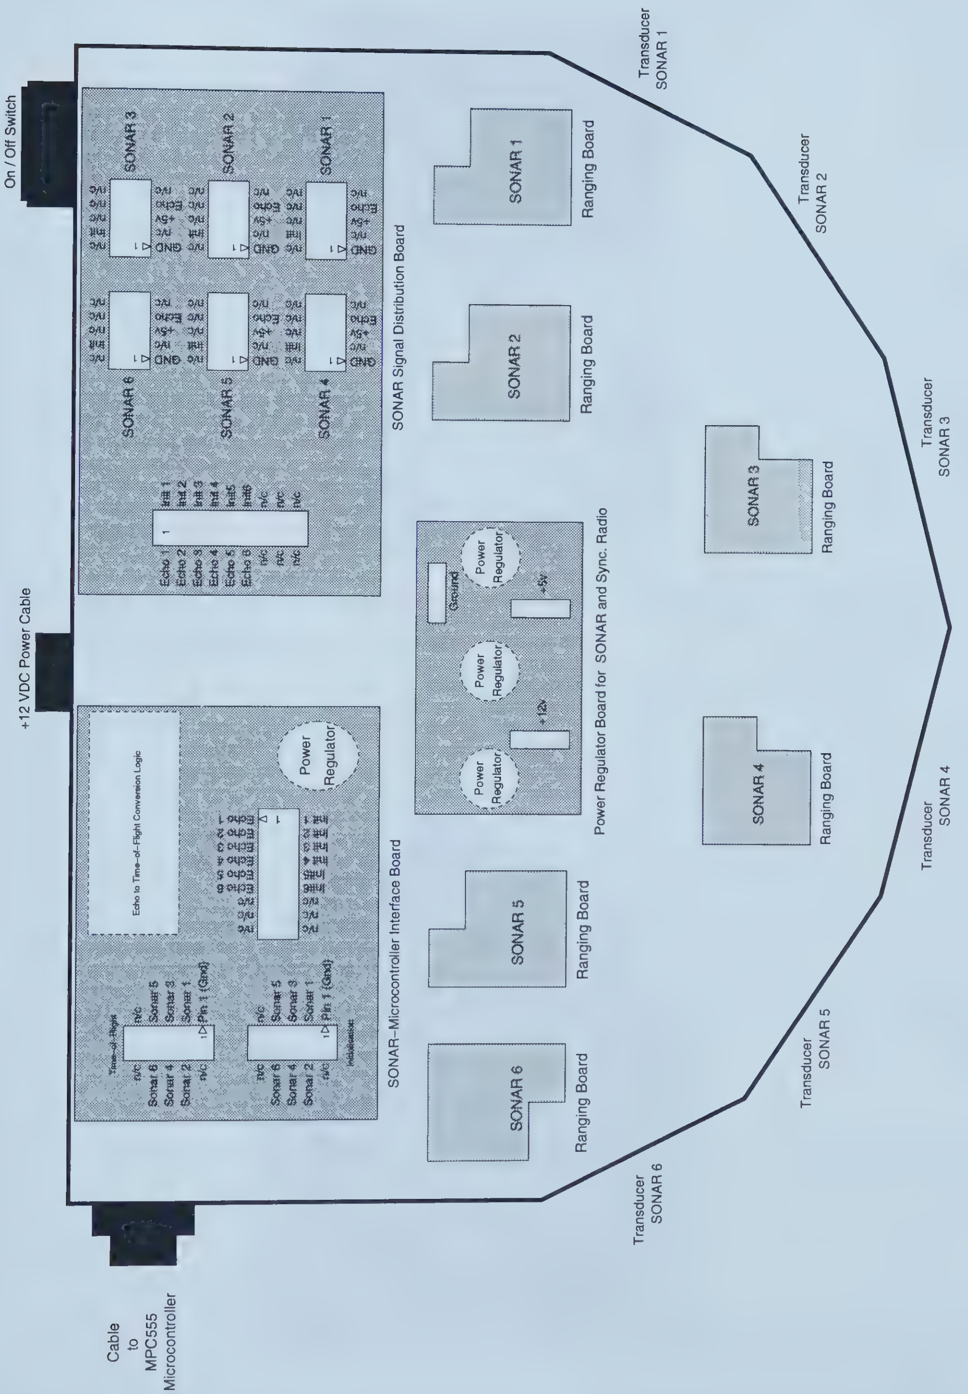

Figure 3.13: SONAR Array board layout

Figure 3.14: Outside SONAR Array. Note the six Polaroid 600 series transducers and the grey cable used for micro-controller signals.

Figure 3.15: Inside SONAR Array. Note the radio synchronisation electronics attached to the Array's lid while the SONAR array electronics are at the bottom.

Figure 3.16: The Complete Pursuit Pack. Note the three Polaroid 9000 series transducers in the centre and the Freewave DGR-115 frequency hopping spread-spectrum radio in the pouch on the right.

Obstacle Detection using the SONAR Array

The SONAR array, shown in Figures 3.13, 3.14, and 3.15, consists of a Lexan polycarbonate resin box housing six multiplexed Polaroid 600 series ultrasonic SONAR transducers, their associated driver circuitry (6500 series SONAR ranging modules) and interfacing circuitry to the micro-controller.

The 600 series electrostatic transducers emit sound in a multi-lobed propagation pattern described in [18, page 2] and are approximated by a 12° symmetrical cone.[8] The six SONAR transducers in the array provide the robot with object detection up to 10 meters away with a field of view of approximately 120°.

3.2.4 Combining Electrostatic and Piezoelectric SONAR Transducers

The three 9000 series piezoelectric transducers used in the Pursuit Pack are designed to meet Society of Automotive Engineers (SAE) specification JI455 for automotive applications [20][9], and emit sound in a multi-lobed pattern [20, page 3] approximated by a 17° by 35° asymmetrical cone, according to [20].[10]

Tests conducted in July 2000 demonstrated that the 9000 series transducer is more capable than the 600 series at detecting obliquely transmitted SONAR pulses (i.e. the receiver is rotated slightly away from the incoming sound). The guard housing on the 600 series seems to prevent it from reliably detecting signals when it is rotated more than 10° from the axis of sound propagation, confirming Godin's findings in [12, page 22]. The 9000 series, with its unobstructed concave transducer can be rotated up to approximately 60° along the axis of its shallowest curve and will still be able to receive SONAR pulses.

Another round of tests demonstrated that while 9000 series transducers, driven by 6500 series ranging boards (with resonance and impedance modifications found in [21] and transmission disabling modifications found in [12]), are excellent receivers of the 50 kHz signals generated by 600 series transducers (driven by unmodified 6500 series boards), they are not efficient transmitters. Unlike the 600 series, which generate an audible clicking sound (in addition to the 50 kHz ultrasonic pulses) each time they are fired, the sound generated by

[8]This approximation is incorrectly stated to be a 12° by 17° asymmetric cone in [19].

[9]One automotive manufacturer now advertises a second benefit to their ultrasonic "Park Distance Control System": it apparently warns hedgehogs and other wildlife of the approaching car.

[10]This is incorrectly stated as 15° by 40° in [19].

the 9000 series transducers has nearly no audible component and the ultrasonic pulses reflects poorly off many surfaces with normally good specular reflection characteristics for ultrasonic signals such as plaster walls or wooden doors. This is likely due to the poor impedance match between the piezoelectric transducer and air as compared to the electrostatic 600 series transducer, yielding a sound pulse of diminished strength. The observation about the 9000 series' poor reflection characteristics was confirmed by a Real World Interface (now a product line of iRobot Corporation) engineer during a conversation with the author at Unmanned Systems 2000 in Orlando, Florida, USA.

For the 9000 series, Polaroid specifies both resonant (transmit) and anti-resonant (receive) frequencies in a band around 45 kHz [20, page 1]. At 50 kHz, the receiving sensitivity of the transducer drops to -70 dB from -65 dB at 45 kHz. Experimentally, it has been shown that this sensitivity is sufficient to detect a 50 kHz signal from a 600 series transducer at a distance of over 10 meters. Therefore, the 9000 series piezoelectric transducer and modified 6500 ranging boards are usable as receivers for the ultrasonic pulses transmitted by the 600 series electrostatic transducers and unmodified 6500 ranging boards.

Because of the 9000 series transducer's compatibility with the 600 series transducers, its identical micro-controller interface due to the use of the 6500 ranging board, its wide field-of-view and its ruggedness it is a suitable candidate for use in this project from the perspective of the designer and the human operator. From the operator's perspective, the ruggedness means that the transducers will last longer in the field. The wider field-of-view will reduce the bulk of the Pursuit Pack because a smaller number of transducers are required.

Figure 3.17 illustrates the connections between the components in the Pursuit Pack. A Motorola 68HC11 micro-controller is responsible for control of the Freewave DGR-115 and synchronisation radios as well as the three 9000 series transducers. The power regulator board ensures that enough electrical power is available to each 6500 series ranging boards; each board can draw two amperes during the momentary initialisation of the SONAR transducer. The SONAR-micro-controller interface board routes Initialisation and Time-of-Flight signals between the signal distribution board and the micro-controller. The signal distribution board routes these signals to and from each of the ranging boards.

The portable Pursuit Pack contains all the necessary components for doing beaconing

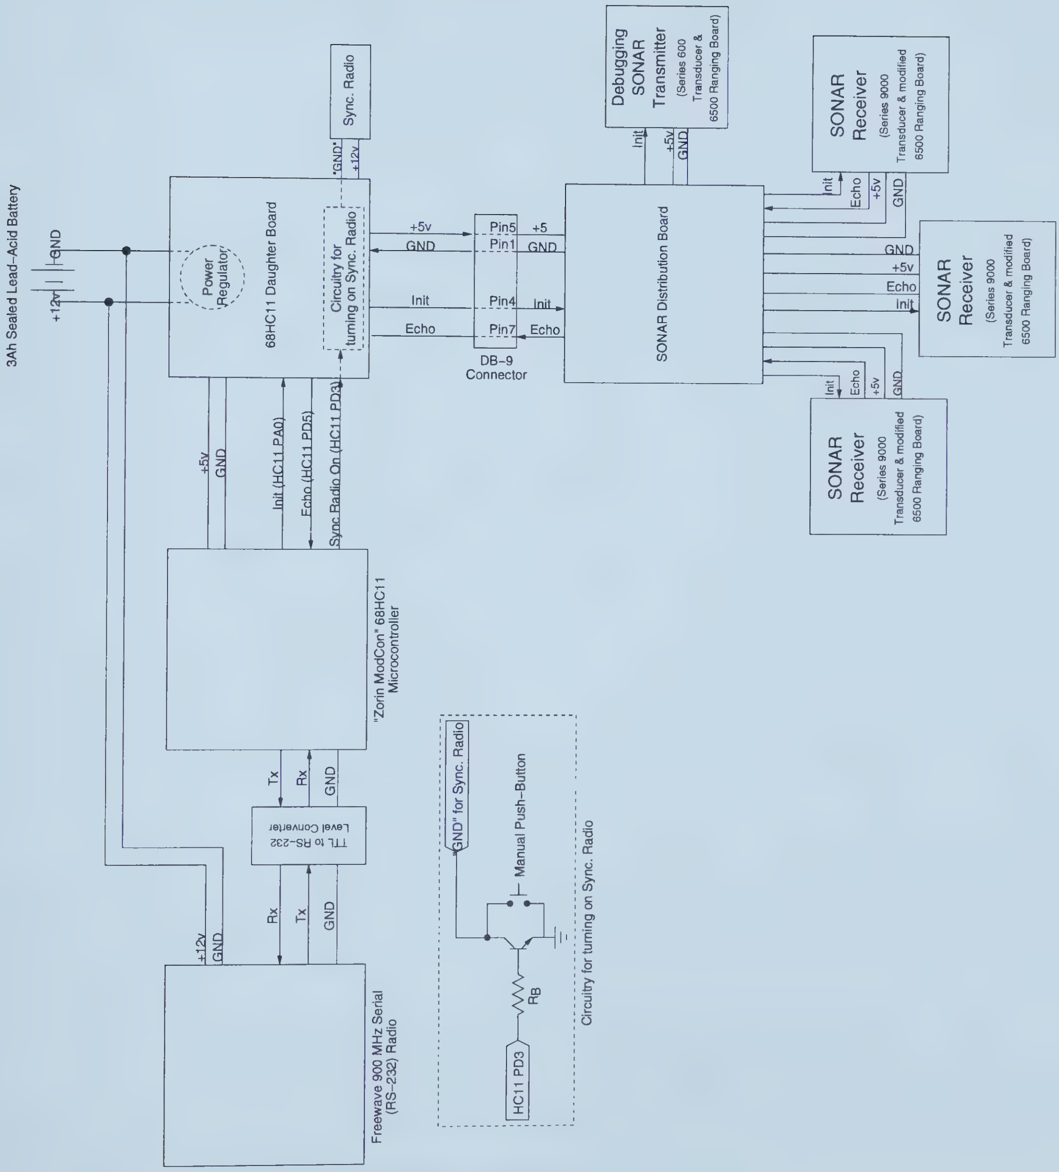

Figure 3.17: The Pursuit Pack Component Block Diagram.

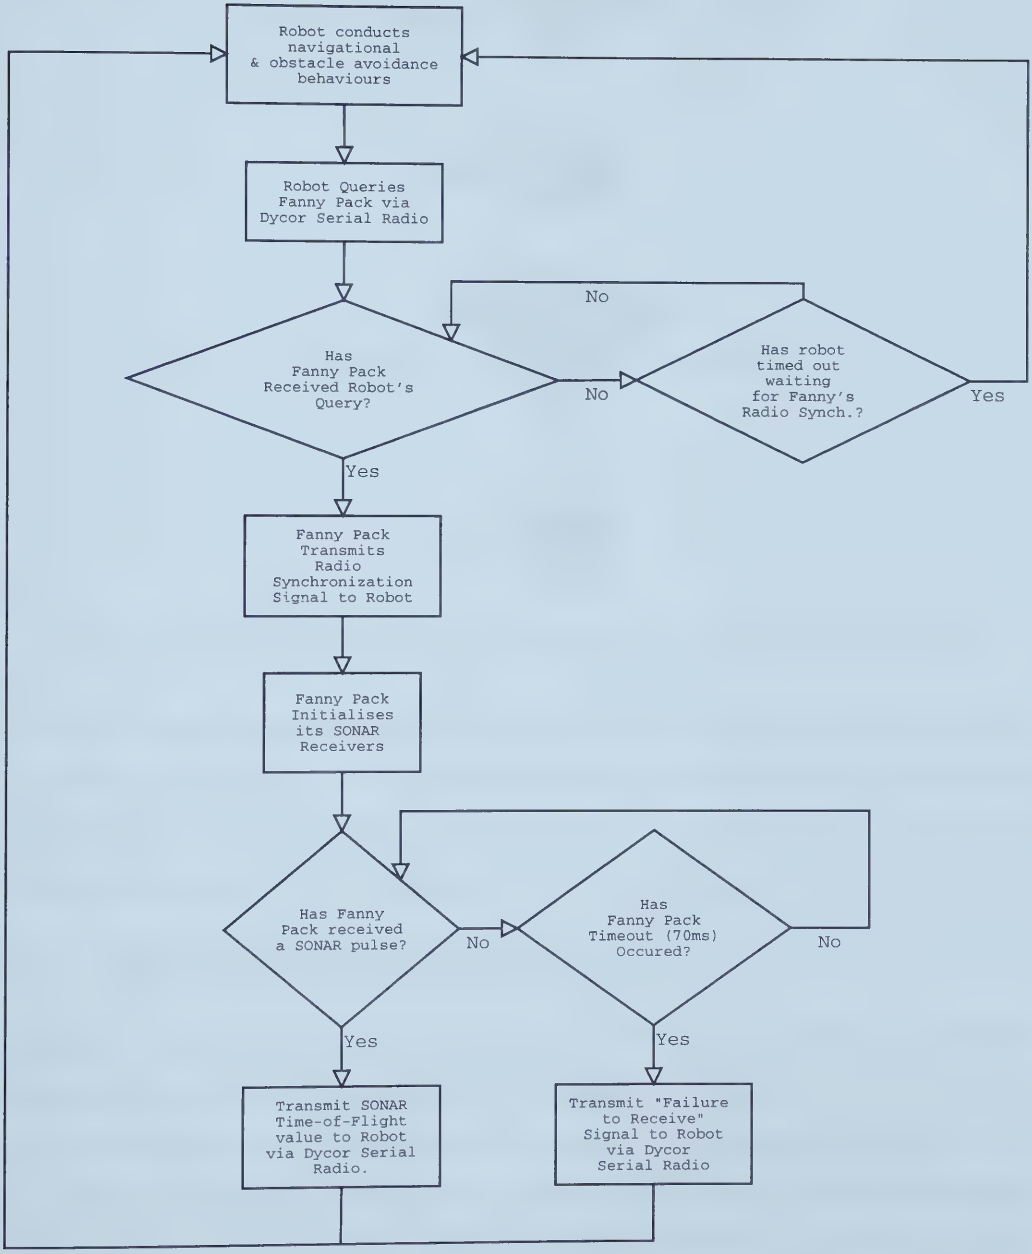

Figure 3.18: Flowchart of Pursuit Pack's high-level operation. Note that Pursuit Pack is referred to as "Fanny Pack."

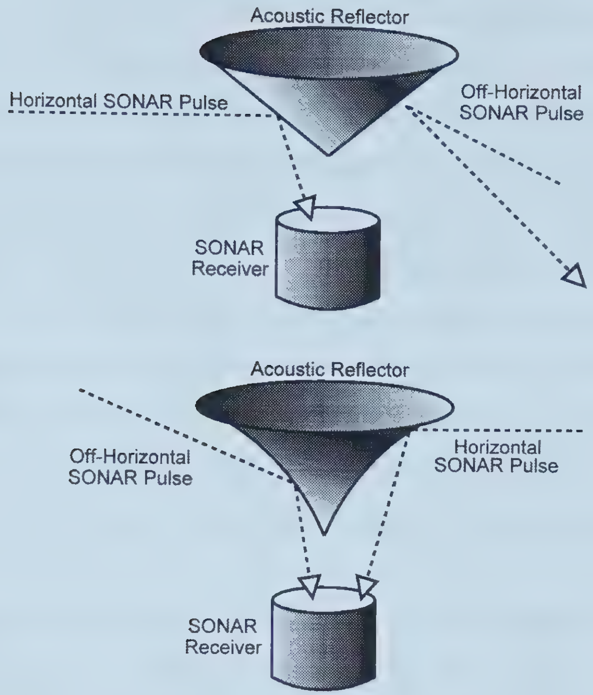

Figure 3.19: A comparison of straight-sloped and convex-shaped acoustic reflectors.

work in an outdoor environment. In the following section the use of the Pursuit Pack as a component in a beaconing application will be discussed. The information gathered by the Pursuit Pack allows the Polar Bear robot to determine distance between the robot and the Pack, an essential task for the implementation of tracking and pursuit.

3.2.5 Pursuit Pack SONAR Beaconing

Although the approach of using trilateration and a reflective cone mounted above a SONAR transceiver (as seen in [9] [5] and Figure 1.1), would seem to be promising for use in this application, it suffers from at least one serious drawback. As can be seen in Figure 3.19 the beacon's acoustic reflector will only allow sound waves travelling in a horizontal direction to be received by the robot. Tests on an aluminium cone with 45° slope, and a maximum diameter of 15 cm. confirmed this to be the case with both the Polaroid 9000 series and 600 series transducers, acting as both transmitters and receivers.

Further tests on a plastic cone with a concave surface, a maximum diameter of 15 cm. and maximum height of 10 cm. demonstrated an improvement in the reception of SONAR

pulses from various non-horizontal directions, as shown in Figure 3.19. This new cone allowed signals of up to 15 ° from the horizontal to be picked up by the receiver at distances of approximately three meters. Three surfaces for this convex cone were experimented with: the original ridged surface, electrical tape and aluminium foil. The ridges covering the cone are a byproduct of the rapid prototyping process used to create the cone from a computer-aided design model. This surface attenuated the signal either because the ridges scattered the sound or the plastic material partially absorbed it. The other two surfaces did not attenuate the signal in any noticeable manner. In addition, it was later realised that the convex curve on this cone contains a discontinuity, resulting in a discontinuity in the possible SONAR transmitter angles. Although this is an improvement over the previous reflector, in field trials this would be completely inadequate if the terrain were even slightly uneven.

The convex acoustic reflector cone did demonstrate improvements over the straight-edged cone. Further research could be conducted into variations in the convex curve, making sure to eliminate any discontinuities like the one observed in this model. As well, two such cones could be used to detect a wider range of angles. This could be accomplished by mounting one on top of the other, with one pointing upwards while the other points downwards. Two SONAR receivers could then be mounted and would be able to detect sound coming from angles above or below the horizontal plane.

To trilaterate position of the beacon in a fashion similar to [9] or [11, pages 425-27] the receiver must be essentially omnidirectional – at least in a plane parallel to the ground. In the context of this project there are at least three possible solutions:

1. A ring of transmitting SONARs worn about the operator's waist.

2. An omni-directional transmitting SONAR beacon placed on the operator's helmet.

3. An omni-directional transmitting SONAR beacon mounted on a pole carried by the operator.

None of these options are ideal. The first solution would require perhaps a dozen or more SONAR transducers and associated support electronics. This would be bulky for the operator to wear. The second solution would probably violate safety helmet standards. For

the last suggested solution a pole would need to be constantly held, reducing the flexibility of the operator. All of these solutions assume that the SONAR beacon worn or carried by the operator transmits SONAR pulses, while the robot receives them. Since the robot's SONAR array and the SONAR transducers to be used with the human operator all function at about 50 kHz the robot's SONAR array cannot be used while range to the operator is being determined. During that time the robot is blind to potential obstacles that it would otherwise be able to detect with its SONAR array.

By not requiring the SONAR system to trilaterate position another option becomes possible. If the SONAR system is only required to determine range and the bearing problem is left to the robot's video system, the equipment (i.e. the Pursuit Pack) worn by the operator no longer needs to transmit sound. Instead, the Pursuit Pack can be used to listen for the incoming SONAR pulses, which are already produced by the SONAR array on the robot.

In the end, the Pursuit Pack, a commercial fanny pack containing three Polaroid 9000 series SONAR transducers, a micro-controller and two radio systems is the most feasible solution in terms of bulk and weight. The three Polaroid 9000 series transducers and accompanying modified 6500 ranging boards, described in Section 3.2.5, are capable of detecting the SONAR pulses transmitted by the 600 series transducer SONAR array on the Polar Bear robot. All six transducers in the array are fired at the same instant and the resulting sound can be detected by the Pursuit Pack's 9000 series transducers over 10 meters away. As shown in Figure 3.21, the three transducers are mounted to maximise the angle at which the Pursuit Pack can receive the robot's SONAR signals. Bearing to the person is determined using the vision system described in Section 3.2.2.

3.2.6 Radio Communication Emphasising Synchronisation

Two radio systems are used for communication between the Pursuit Pack and the Polar Bear robot: a pair of spread-spectrum frequency hopping serial transceivers and a single frequency transmitter and receiver pair. The Freewave DGR-115 spread-spectrum radio transceivers are used for the transmission of status and telemetry information. The robot's on-board PC relays information from the robot's micro-controller to one Freewave transceiver via on of its two serial ports while the other transceiver is directly attached to the Pursuit Pack micro-controller's serial port. Because the variable latency in the packet-based

signal transmission inherent in this type of radio (as well as wireless local area network radios and similar radios discussed on page 9) negatively affects the process of measuring the time-of-flight of SONAR pulses a fixed latency transmitter and receiver pair is required for synchronisation. A single frequency transmitter and receiver pair salvaged from an automotive alarm system serves this purpose.

The receiver and transmitter radio pair used for synchronisation have a fixed signal latency of 75 ms. A circuit board, shown in Figure D.7, eliminates pulsed noise in the signal and converts the signal to Complimentary Metal Oxide Semiconductor (CMOS)-compatible logic levels. The transmitter is directly controlled by the Pursuit Pack 68HC11 micro-controller while the receiver is directly monitored by the robot's MPC555 micro-controller.

The resulting radio system allows relatively complex information to be transmitted via the spread spectrum radios while the simple synchronisation signal is communicated by a single frequency transmitter-receiver pair. The process begins with the robot querying the Pursuit Pack via the spread spectrum radio prior to a beaconing attempt. The Pursuit Pack's 68HC11 micro-controller acknowledges the robot's query via the synchronisation radio.

Because it was originally used as an automotive car alarm control the synchronisation radio transmitter actually transmits a series of pulses which match an identical series known by a micro-controller embedded in the receiver. This series of pulses is obtained from the car alarm receiver board and modified by the board shown at the top of Figure D.7. Specifically, the operational amplifier and a filtering capacitor convert the received sinusoidal pulse train into a single on/off pulse. The CMOS nand gate is used to ensure that the signal has only +12 volt or 0 volt values (i.e. the power rails of the radio receiver). An opto-isolator and another CMOS nand gate are then used to convert the signal to +5 volt / 0 volt levels which can be measured by the robot's MPC555 micro-controller.

Once the MPC555 micro-controller receives the synchronisation signal it commands the six ranging boards in the robot's SONAR array to release a single burst of ultrasonic pulses. The time taken by the first SONAR pulse to reach the Pursuit Pack is recorded and sent back to the robot via the spread spectrum radio. (Of course, the 75 ms. latency in the synchronisation radio signal is taken into account by the Pursuit Pack's micro-controller when calculating the time-of-flight.) If no SONAR pulse is received by the Pursuit Pack

before an adjustable time-out the robot is alerted via the spread spectrum radio. The sequence of these radio and SONAR transmissions is illustrated in Figure 3.20.

3.2.7 Combining SONAR Beacon and Vision Information

Instead of trilaterating the position of the operator using a single SONAR transmitter and a minimum of three receivers (as is done in [11, pgs 426-427] and [9]), this system combines the SONAR range information with the vision's bearing information. Because the shape and size of the operator's vest changes, reliable range information cannot be obtained from the camera. Likewise, without multiple SONAR receivers mounted on the robot and sufficiently spaced apart bearing information is not possible using time-of-flight measurements.

Speed of the vehicle is based on the distance between the robot and the operator, as reported by the Pursuit Pack. Fully adjustable, the robot will speed up if the operator is more than three meters away and will slow down and stop if the operator is less than one meter away. Direction of the robot is based on the location of the operator's safety vest. The robot attempts to turn so that the vest is centred in the camera image: if the vest is in the right-hand part of the image the robot will turn to the right, and it will turn left if the vest is seen in the left-hand part of the image.

If communication with the Pursuit Pack is severed or if it reports that it cannot receive the robot's SONAR pulses the robot will come to a complete stop and wait until this condition changes. Likewise, if the robot loses sight of the operator it will come to a complete stop.

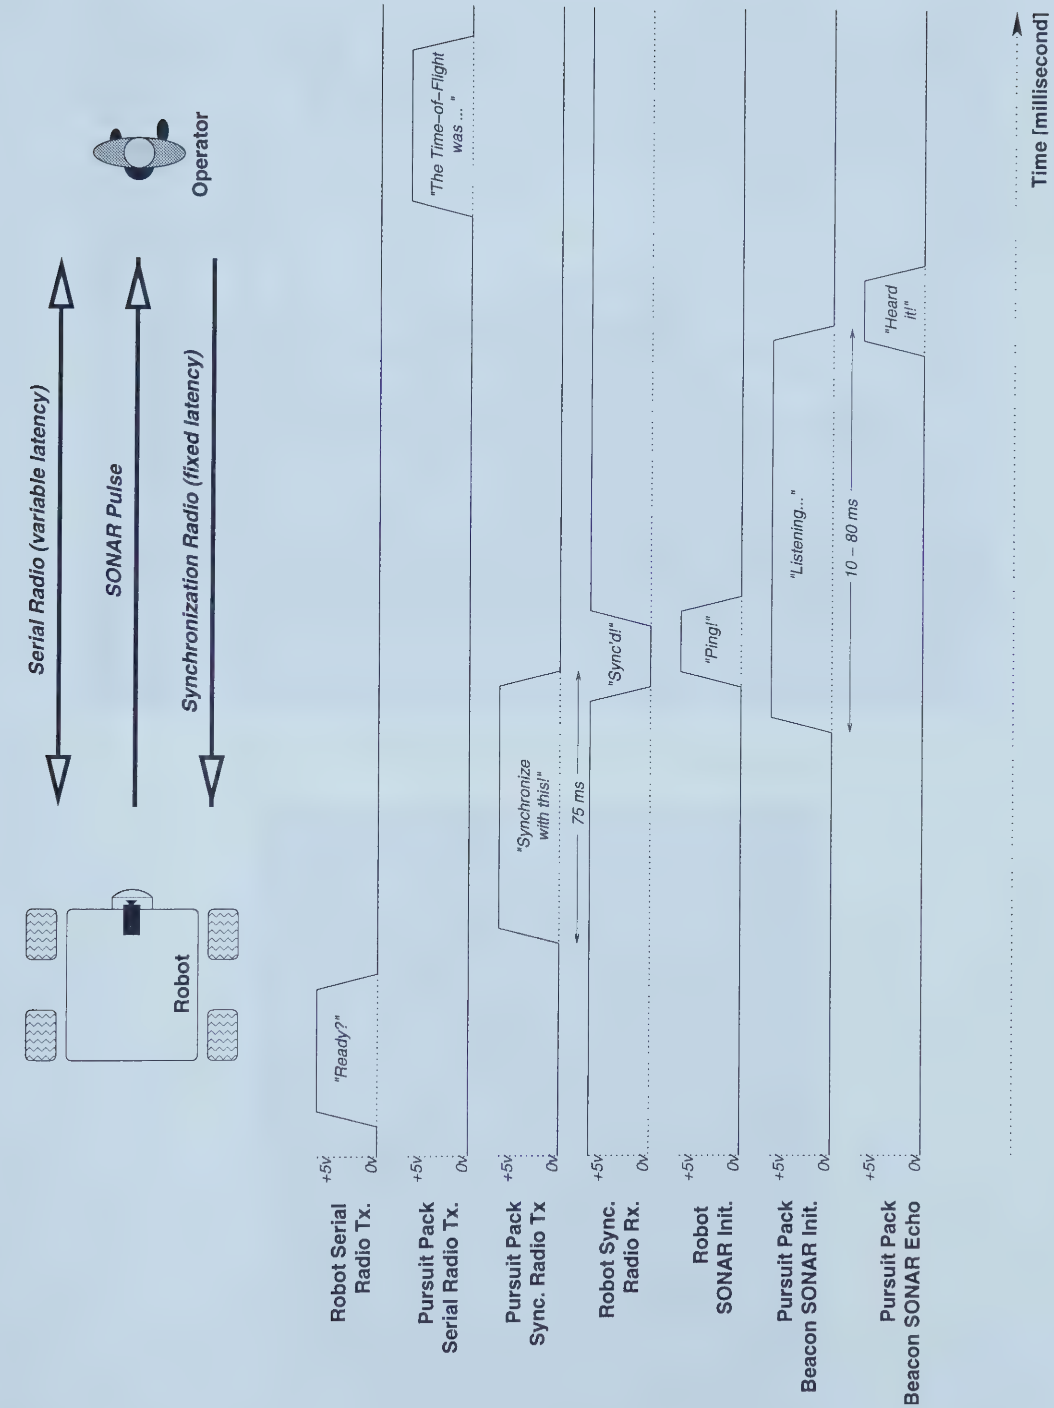

Figure 3.20: The communication signals between the Polar Bear and the Pursuit Pack. Note: this is not to scale. Length of signals created by the Freewave serial radios is variable; only the rising edges of the SONAR signals are important, so duration has not been measured.

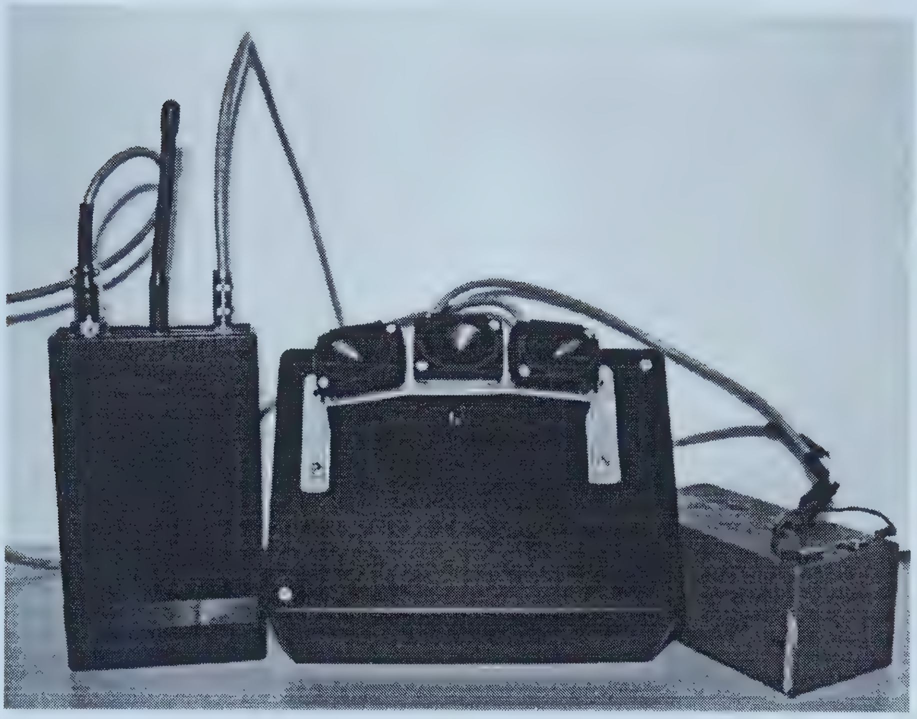

Figure 3.21: The Pursuit Pack removed from fanny pack.

Figure 3.22: The Pursuit Pack, worn.

Chapter 4

Experimental Results

4.1 The Experimental Setup

Three technology demonstrations directly related to the SAGV project have been held: Calgary, Alberta in January 1999; Rocky Mountain House, Alberta in August 1999; and Edmonton, Alberta in September 2000. All three demonstrations are discussed below, with particular emphasis placed on the final demonstration held in September 2000.

4.1.1 January 1999: Calgary, Alberta

This demonstration was held in an empty lot next to the Schlumberger building in downtown Calgary. The Polar Bear robot was tele-operated using an on-board 68HC11 micro-controller, an off-board laptop computer and a pair of Freewave DGR-115 radios connecting the two systems. Direction commands were typed into the laptop by a person and these commands were then issued to the robot via the two radios. The robot was not equipped with any sensors.

The environment was relatively flat and was composed of compacted dirt, asphalt and concrete.

4.1.2 August 1999: Rocky Mountain House, Alberta

The second demonstration occurred in a gravel pit outside Rocky Mountain House, Alberta in August 1999. Again, the robot was tele-operated using a pair of Freewave DGR-115 radios connecting the robot's on-board 68HC11 micro-controller to an off-board laptop computer. The robot was equipped with an array of six SONAR transducers for detecting

object range and movement and the laptop ran a real-time program which analysed the SONAR sensor data and returned navigational commands to the robot.

The gravel pit's uneven and bumpy floor was a mixture of gravel, wet mud and dry mud, with loose gravel covering most of the walls of the pit. A large standing pond with a maximum depth of half a meter was available in the gravel pit.

4.1.3 September 2000: University of Alberta, Edmonton, Alberta

On September 25 2000 the latest technology demonstration was held outside of the Mechanical Engineering Building on the main University of Alberta campus. The robot housed a MPC555 micro-controller and an on-board battery-powered personal computer. A Sony Handycam video camera and six SONAR transducers provided visual and range-finding sensing capability. The robot could now be led by a person equipped with a waist-mounted SONAR beaconing package and a standard orange safety vest.

The area on the north side of the Mechanical Engineering Building contains mild, grassy terrain, a number of large trees, a concrete sidewalk, and a flat asphalt road lined with round concrete curbs.

General Demonstration Results			
Category	*Calgary, 1999*	*Rocky Mtn House, 1999*	*Edmonton, 2000*
Duration	10 min.	90 min.	25 min.
Distance	50 m.	400 m.	200 m.
Max Speed	18 km/h	8 km/h	5 km/h
Electrical Life	indefinite	indefinite	16 hrs (est)
Turn Radius	0 m.	1.5 m.	3 m.
Terrain Type (1)	flat	moderately flat, broken	flat
Terrain Type (2)	asphalt	gravel	asphalt, grass
Obstacles Present	no	no	sparse
Obstacle Type	n/a	n/a	large trees
Water	no	yes	no
Water Depth	n/a	50 cm.	n/a

4.2 General Polar Bear Results

The maximum speed at which the robot was tele-operated during the Calgary demonstration was about 15 km/h (equivalent to a moderate jog). During the demonstration in Rocky

Figure 4.1: Polar Bear pursuing Jim Qualie. (University of Alberta, Edmonton, AB; 2000)

Mountain House the robot achieved a maximum speed of about 8 km/h (a fast walk), while the maximum speed for the robot during the Edmonton demonstration was approximately 5 km/h (a moderate walk). The robot only reverses when it is being manually tele-operated.

The reliability of the robot in pursuing a person was much higher during the Edmonton demonstration than during the Rocky Mountain House one. The robot uses its SONAR array far more effectively and will now reliably (i.e. it did not fail once during tests) stop if the operator, or any other large obstacle, is one meter or less from the robot's array. (This stopping distance is reprogrammable.)

The robot can operate in all of the environments described above, including the standing pond found in the gravel pit outside Rocky Mountain House. Because feedback from the wheels or from the hydraulic system is not measured the robot cannot traverse loose soil or gravel if the terrain is inclined more than a few degrees.

Note that the increase in turning radius is due to limitations in the robot's field of view. In Calgary the robot was tele-operated and did not have a field-of-view limitation. In Rocky Mountain House the robot would lose track of the person it was following if the person

attempted to make the robot turn in a circle with a radius smaller than 1.5 meters. Likewise, in Edmonton, where the limitation was due primarily to the loss of visual identification when the person stepped outside of the camera's field of view.

4.2.1 Deploying Geophysical Cable

The Polar Bear showed in both the Rocky Mountain House and Edmonton demonstrations that it is capable of automatically deploying geophysical cable at a rate approximately equal to the speed of the robot.

4.3 SONAR Array Results

The SONAR array, used to detect objects in front of the robot, has been improved between the Calgary and Edmonton demonstrations. The first array, with two forward-looking transducers and four angled ones, produced a 90 ° ultrasonic field with two major blind spots. Also, the maximum detectable range was only slightly more than three meters due to a limitation in the software routines used to control the micro-controller's internal timer (which measured the time-of-flight of the SONAR pulses).

The six transducer SONAR array used in the Edmonton demonstration produces a 120 ° field with no observable blind spots. The maximum range (for detecting an averaged sized adult) when the array is used for regular range-finding (i.e. not communicating with the Pursuit Pack) is just slightly less than the 10.7 meters specified by Polaroid Corporation.

4.4 The Pursuit Pack & Associated Vision System Results

As can be seen in Figures 3.12 and 4.1 a tether connects the Pursuit Pack to the Polar Bear robot. This tether bypassed the synchronisation radio transmitter and receiver during the technology demonstration on September 25, 2000. During indoor trials in the week prior to the technology demonstration the synchronisation radio system successfully linked the Pursuit Pack and the Polar Bear (in the method described in Section 3.2.6) at distances of up to approximately 10 meters (limited by the confines of the room) and aided the complete system to yield accurate range data with a maximum 5% uncertainty. In trials two days prior to the demonstration the signal produced by the synchronisation radio receiver began to

deteriorate to the point of being unreliable. The reason for this deterioration is unknown but could be related to the method in which the transmitter is turned on and off (by a transistor which alternatively grounds the transmitter. See Figure 3.17.). The time constraint required an immediate remedy and, because attempts to repair the synchronisation radio failed, a tether connecting the RS-232 port on the robot's PC to the RS-232 port on the Pursuit Pack's micro-controller was made. This tether, similar in concept to the one described in [11, page 426], also rendered the Freewave spread-spectrum radio redundant because it was able to conduct fixed-latency serial transmissions between the Polar Bear robot and the Pursuit Pack.

During the demonstration the Pursuit Pack's SONAR array was only two-thirds operational. One of the 9000 series transducers was disconnected due to anomalous output prior to the demonstration. The removal of the transducer in no way affected the performance of the robot during the demonstration; during the instances in which the robot failed to detect Jim Qualie, the operator, the cause of the failure was not related to the disconnected transducer but was rather due to extreme lighting conditions which interfered with the vision processing responsible for identifying Jim's orange safety vest.

In some instances when the operator would step out of the shade and into bright sunlight the robot's vision system would lose sight of the orange vest because it had become so bright. The robot would come to a halt and attempt to reacquire the vest; occasionally the vest's colour would have to be manually reset via the on-board PC because the robot could not do so on its own. As well, just like a human driver, the robot would sometimes become blinded when driving towards the sun and the operator attempted to shade the camera lens in order to have the robot continue.

The combination of vision and SONAR for pursuing a human operator proved to be very effective, blinding of the vision system by the direct sunlight notwithstanding. The maximum distance from the robot that the operator can be, for the robot to follow effectively, is between three and four meters. This limitation is due to the size of the safety vest worn by operator. A higher resolution camera or a single coloured coverall suit worn by the operator could increase the effective range.

Vision-only and SONAR-Vision Tracking Results		
Category	*Vision-Only (P-3-P)*	*SONAR-Vision*
Max Roll (Visual Target)	20 °	360 °
Roll Uncertainty	5 °	n/a
Max Pitch (Visual Target)	20 °	not measured
Pitch Uncertainty	5 °	n/a
Max Yaw (Visual Target)	60 °	not measured
Yaw Uncertainty	5 °	n/a
Max Roll (Pursuit Pack)	n/a	360 °
Max Pitch (P. Pack)	n/a	17 ° (est.)
Max Yaw (P. Pack)	n/a	70 ° (est.)
Range: Visual Target	1.5 - 7 m.	7 m. (max), 3 m. (usual)
Range Error	5%	n/a
Range: SONAR Obstacle	n/a	10 m.
Range Error	n/a	2%
Range: SONAR Beacon	n/a	10 m. (min)
Range Error	n/a	5%
Field of View: Vision	90 °	90 °
Field of View: SONAR	n/a	120 °

4.5 Vision-only Tracking System Results

This system has only been tested in two indoor environments, the Advanced Robotics and Tele-operation Lab and the Mechanical Engineering Machine Shop at the University of Alberta. In both cases, the ambient light was mostly artificial.

The visual tracking system discussed in Section 3.1 is able to determine the pose of the target with surprising accuracy considering the resolution (320 by 200 pixels) of the image and the use of a commercial camcorder. Experimentally, the maximum X-axis (pitch) rotation that can be detected is about $\pm 20^\circ$, the maximum Y-axis (yaw) rotation is about $\pm 60^\circ$, while the maximum Z-axis (roll) rotation is about $\pm 20^\circ$. All rotations are measured to one degree resolution with an accuracy of approximately $\pm 5^\circ$.

The usable range for target distance along the camera's optical axis has been found, experimentally, to be 1.5 meters to 7 meters $\pm 5\%$. System response is about 4.25 frames per second with a delay of about one or two frames due to the image buffer in the capture card. The frame rate could be increased to between seven and ten if the portions of the programming code related to inter-computer communication (useful for tele-operation) were removed. This increase in frame rate was demonstrated by Jim Qualie in the code

optimisation he performed for the vision component described in Section 3.2.2.

A visual representation of the resulting system can be seen in Figure 3.9 found on page 48.

Chapter 5

Future Work and Conclusions

5.1 Future Directions for Work

This project continues to be an exciting one in which to conduct research because it offers a great deal of hands-on development not normally available to graduate students.

Although there has been much success in the development of a robot capable of following a person in outdoor environments, there is still much work to be done before it can be used in the real world. Most of the work involves improving the sensors on the robot. Here are a list of suggestions for possible future work:

1. Modify the 600 series transducers to emit at more than 50 kHz
   - Crosstalk between obstacle detection SONARs and beacon detection SONARs can be avoided if the two systems used separate frequencies (one at 45/50 kHz while the other is at a higher frequency)
   - Purchase Polaroid's Ultrasonic Developers Kit to conduct initial research into alternative frequencies
2. Investigate SICK Corporation LASER range-finder ability to detect reflective material
   - Some SICK range-finders have the ability to identify reflective material such as those found on safety vests
   - This property is not widely advertised but is mentioned in the SICK data sheets and SICK engineers are familiar with it

3. Investigate frequency-sensitive SONAR for beaconing
   - Australia's Hexamite.com sells a variety of alternative ultrasonic beaconing systems which could be useful
4. Addition of a second video camera
   - To complement the current one by providing a greater field of view or over-all resolution
   - This is already being seriously considered by the ARVP student group
5. Addition of additional safety components
   - Infrared range-finders (0.8 m. maximum range) along the perimeter of the chassis (Supplier: HVW Technologies)
   - Additional emergency stop buttons
   - Replacement of the current radio emergency stop
6. Further Development of Vision-only Tracker
   - Integrate it with the SONAR beaconing system and attempt to have the Polar Bear follow a truck equipped with a five feature point target
   - Monitor progress in the Intelligent Ground Vehicle Competition's Follow-the-Leader event since the organisers have adopted a similar target based on the author's recommendations

5.2 Conclusions

5.2.1 The Polar Bear Platform

The Polar Bear is a unique university student built robot. Its rugged mechanical system, expandable electronics and adaptable software make it an ideal platform for outdoor mobile robotics. The electronics and mechanical components have been designed to be reliable and easily serviceable. The experimental trials conducted on the Polar Bear would not have been possible with the vast majority of commercial or research robotics platforms.

Because of a focus on design and construction for further development in industry Polar Bear is beginning to find a niche in the oil industry. The current system is rated at Level 6 (out of a possible 9) on the NASA Technology Readiness Levels scale, meaning that it is relatively mature from a developmental point of view and that commercially realisable goals in the near future are definite possibilities.

5.2.2 The Vision-only Tracking System

A vision system, separate from, but related to, the SONAR beaconing has been introduced in this thesis. This system demonstrates how the Perspective-3-Point machine-vision algorithm can be used in an iterative fashion to solve a real-time visual tracking problem using inexpensive hardware. The status of this project is Level 4 ("Component and/or breadboard validation in laboratory environment") on the NASA Technology Readiness Levels scale. It is currently ready to be integrated with the other systems on the Polar Bear robot and be run through vehicle tracking field trials in a relevant outdoor environment.

5.2.3 Tracking using Vision & SONAR Beaconing

Mixture of Electrostatic and Piezoelectric SONAR Transducers

To the author's best knowledge, no other robotics project has demonstrated a complementary use of electrostatic and piezoelectric SONAR transducers in a manner such as that seen in this thesis. By combining an electrostatic transmitting transducer, with its high signal strength, and a piezoelectric receiving transducer, with its wider field-of-view, a beaconing system capable of being used outdoors has been developed. Although it does require more transducers than a comparable system using a cone-shaped acoustic reflector mounted above a single transducer, it is far more capable in a non-laboratory environment than the latter system.

It should also be noted that the development of the Pursuit Pack is, to the author's best knowledge, the only successful implementation of the Polaroid 9000 series piezoelectric transducer in a mobile robotics project.

Introduction of Radio Synchronisation

Although the concept of radio synchronisation is not new, its application to SONAR beaconing systems in ground robotics seems to be. None of the projects surveyed which concerned SONAR beaconing addressed this issue in a meaningful manner. This thesis successfully illustrates how, in a SONAR beaconing application, two radio systems can be used in a complementary fashion such that one communicates synchronisation information using a method with fixed and known latency, while the other communicates more complex but less time sensitive information using another method which has variable and unknown latency. The latter is necessary because of the data transmission techniques inherent in most packet-based serial data radios, while the former is possible because of the simplistic nature of a single synchronising radio pulse or pulse train.

Augmentation of SONAR with Vision

By adding a vision component to the SONAR array on the robot, the development of the Pursuit Pack beaconing system has been simplified since trilateration is not required. The vision system, although rather simple at this point, is flexible and can be easily improved using techniques common in the machine vision community such as the Perspective-3-Point algorithm described earlier.

This current vision and SONAR system has shown that it is capable of tracking a person in real time and allows the Polar Bear robot platform to pursue the person at speeds of up to 5 km/h.

Bibliography

[1] Margaret Cheney. *Tesla: Man out of Time*. Prentice-Hall, New Jersey, 1981.

[2] Sarah Price. *Pose Estimation from Corresponding Feature Data*. On-Line Compendium of Computer Vision, no date. ONLINE Available: http://www.dai.ed.ac.uk/CVonline/LOCAL_COPIES/MARBLE/high/pose/pose.htm [April 6, 2000].

[3] Hong Zhang. *Perspective-n-point problem – Errata*. University of Alberta, no date. ONLINE Available: http://www.cs.ualberta.ca/~zhang/c631/errata.htm [April 6, 2000].

[4] Long Quan and Zhongdan Lan. Linear n-point camera pose estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8), August 1999.

[5] J. Borenstein, H. R. Everett, and L. Feng. *Navigating Mobile Robots: Sensors and Techniques*. A. K. Peters, Ltd., Wellesley, Massachusetts, cd-rom edition, n/a. Out of print. Files available online at http://www-personal.engin.umich.edu/~johannb/my_CDROM.htm.

[6] Oxford University Press, Oxford, United Kingdom. *Oxford Illustrated Dictionary*, second edition, 1985.

[7] Maritime Faculty, Southampton Institute, Southampton Hampshire, United Kingdom. *Extending Control: Triangulation*. ONLINE Available: http://www.solent.ac.uk/hydrography/notes/horizcon/extend/extend2.htm [October, 2000].

[8] Maritime Faculty, Southampton Institute, Southampton Hampshire, United Kingdom. *Extending Control: Trilateration*. ONLINE Available: http://www.solent.ac.uk/hydrography/notes/horizcon/extend/extend3.htm [October, 2000].

[9] L.E. Navarro-Serment, C.J.J. Paredis, and P.K. Khosla. A beacon system for the localization of distributed robotic teams. In *Proceedings of the International Conference on Field and Service Robotics*, Pittsburgh, PA, USA, August 1999.

[10] Robert Langreth. The bad boy of robotics. *Popular Science*, 246:88–92, 1995.

[11] H.R. Everett. *Sensors for Mobile Robots: Theory and Application*. A.K. Peters, Ltd., Wellesley, Massachusetts, 1995.

[12] Alysha Godin. Implementation of a Multidimensional Ultrasonic Position Measurement System. Thesis supervisor: C.P. Diduch, Department of Electrical and Computer Engineering, University of New Brunswick., April 1999.

[13] Parker Hannifin Corp. *Hydraulic Products and Total Systems Engineering Catalogue*, unknown.

[14] U.K. Deiters. *Subroutines for Solving Cubic, Quartic or Quintic Equations*. ONLINE Available: http://www.uni-koeln.de/math-nat-fak/phchem/deiters/quartic/quartic.c [April 6, 2000].

[15] James Andrew Smith. Computing science 631 class notes. University of Alberta, 2000.

[16] Sony Corporation (Semiconductor Business Division). *ICX208AKB Diagonal 4.5mm (Type 1/4) CCD Image Sensor for NTSC Color Video Cameras*, no date. ONLINE Available: http://www.sel.sony.com/semi/PDF/ICX208AKB.pdf [April 1, 2000].

[17] SICK Optic. *Laser Measurement Systems - OUTDOOR*, no date. ONLINE Available:http://www.sick.de/english/products/dat_products.asp?bannerback=0 [November, 2000].

[18] Polaroid OEM Components Group. *Technical Specifications for 600 Series Instrument Grade Electrostatic Transducers*, 1999. ONLINE Available: http://www.polaroid-oem.com/pdf/electrans.pdf [November, 2000].

[19] Polaroid OEM Components Group. *Polaroid OEM : Ultrasonics*, no date. ONLINE Available: http://www.polaroid-oem.com/ultrason.htm [November, 2000].

[20] Polaroid Corporation. *Technical Specifications for 9000 Series Piezo Transducer*, 1998. ONLINE Available: http://www.polaroid-oem.com/pdf/9000series.pdf.

[21] Acroname Inc. *Driving the Series 9000 Transducer with a 6500 Ranging Board*, 2000. ONLINE Available: http://www.acroname.com/robotics/info/ideas/9000/9000.html [November, 2000].

[22] Schlumberger Ltd. *Annual Report*, 1999. ONLINE Available: http://www.1.slb.com/ir/ar/ar99/slb_org.html [October, 2000].

[23] Schlumberger Geco-Prakla. *Background Information*, no date. ONLINE Available: http://www.geco-prakla.no/systemsider/eBI.htm [October 17, 2000].

[24] Geosphere, Inc. *Seismic Reflection and Refraction Methods*, September 2000. ONLINE Available: http://www.geosphereinc.com/seismic.htm [October 15, 2000].

[25] TerraDat Geophysical - U.K. *Seismic Surveying in General*, no date. ONLINE Available: http://www.terradat.co.uk/seismic1.html [October 15, 2000].

[26] Schlumberger Geco-Prakla. *Navpak: Inertial Navigation System*, no date. Brochure OCG-5509.

[27] University of Alberta. Semi-autonomous guided vehicle video. Ten copies made available to project stake-holders, September 2000.

[28] Armed Forces Journal International. MAGAZINE, July 2000. Available: http://www.afji.com/AFJI/Mags/2000/July/ [November 25, 2000].

[29] The Economist. *Look, No Pilot*, November 9 2000. ONLINE Available: http://www.economist.com/science/displayStory.cfm?Story_ID=417775.

[30] University of Victoria, Department of Mechanical Engineering, Faculty of Engineering. *Space and Subsea Robotics Laboratory [Webpage]*, July 2000. ONLINE Available: http://subspace.me.uvic.ca/ [December 24, 2000].

[31] Unmanned Underwater Vehicle Showcase. *Unmanned Underwater Vehicle Showcase 2000 Southampton [Webpage]*, no date. ONLINE Available: http://www.uuvs.net/ [December 24, 2000].

[32] Naval Meteorology and Oceanography Command [of the United States]. *Unmanned underwater vehicle demonstration gives Navy a glimpse of the future in ocean surveying*, December 22 1998. News Release No. 98-032 ONLINE Available: http://pao.cnmoc.navy.mil/PAO/News/Press%20Rel/Pres1998/32-UUV%20Fest.htm [December 24, 2000].

[33] Hoa G. Nguyen and John P. Bott. Robotics for Law Enforcement: Applications Beyond Explosive Ordnance Disposal. In *SPIE International Symposium on Law Enforcement Technologies*, Boston, MA, USA, November 2000. Available: http://www.spawar.navy.mil/robots/pubs/SPIE4232A-69.pdf [November, 2000].

[34] John Blitch. *Tactical Mobile Robots*. Defense Advanced Research Projects Agency, no date. ONLINE. Available: http://www.darpa.mil/ato/programs/tmr.htm [November, 2000].

[35] Sandia National Laboratories. *Robotics and Intelligent Machines in the U.S. Department of Energy; A Critical Technology Roadmap, Executive Summary*, October 1998. SAND98-2401/1.

[36] S. Baten. Autonomous Road and Contour Following with a Tracked Vehicle. In *Unmanned Ground Vehicle Technology II*, volume 4024. SPIE 14th Annual International Symposium on Aerospace/Defence Sensing, Simulation, and Controls, April 2000.

[37] I. Schwartz. PRIMUS: Autonomous Driving Robot for Military Applications. In *Unmanned Ground Vehicle Technology II*, volume 4024. SPIE 14th Annual International Symposium on Aerospace/Defence Sensing, Simulation, and Controls, April 2000.

[38] G. Garrison. *Surveillance and Reconnaissance Ground Equipment (SARGE)*. Intelligent Systems and Robotics Center, Sandia National Laboratories, September 2000. Available: http://www.sandia.gov/isrc/Capabilities/Integration_Technologies/SARGE/sarge.html [November, 2000].

[39] Keith Miller. *DIXIE Surveillance Robot*. Intelligent Systems and Robotics Center, Sandia National Laboratories, November 1999. ONLINE Available: http://www.sandia.gov/isrc/Capabilities/Integration_Technologies/DIXIE/dixie.html [November, 2000].

[40] Robotics Institute, Carnegie Mellon University. *[Transitional Unmanned Ground Vehicle] Project Homepage*, no date. ONLINE Available: http://www.frc.ri.cmu.edu/~ssingh/tugv.html [September, 2000].

[41] RIEGL Laser Measurement Systems GmbH. *Laser Rangefinder on a Robot System in Antarctica*, March 2000. ONLINE Available: http://www.riegl.co.at/an_hs027.htm; RIEGL HS027 (3/99).

[42] David Kortenkamp, R. Peter Bonasso, and Robin Murphy, editors. *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*. American Association for Artificial Intelligence Press / The Massachusetts Institute of Technology Press, Cambridge, Massachusetts, 1998.

[43] Unmanned Systems 2000 – Plenary Speaker Biographies. Symposium Handout, Orlando, Florida, July 2000.

[44] Mike Toscano. New Developments in the [Office of the Under Secretary of Defense] Joint Robotics Program (JRP). In *Unmanned Ground Vehicle Technology II*, volume 4024. SPIE 14th Annual International Symposium on Aerospace/Defence Sensing, Simulation, and Controls, April 2000.

[45] Aeronautical Systems. *New Unmanned Earth Science Platform to be Developed Under NASA ERAST Program*, January 2000. General Atomics News Release.

[46] John Newton. Autonomous and remotely guided vehicle market study. Technical report, John Newton & Associates, 168 Willow Way, Edmonton, AB, T5T 1C8, July 2000. Prepared for Western Economic Diversification Canada; Available: http://www.ee.ualberta.ca/ arvp/docs_2000/ARGVMarketStudy.pdf [November, 2000].

[47] Randy Frank, Neil Krohn, and Herb Shukman. *The Role of Semiconductor Sensors in Automotive Power Train and Engine Control*. Sensors Magazine, no date. ONLINE Available: http://www.sensorsmag.com/articles/1298/rol1298/main.shtml [October 30, 2000].

Appendix A

The Geophysical Industry

A.1 State of the Art

Oilfield Services is one of the three reportable business segments of Schlumberger Limited. It provides services to its clients involved in the exploration and production of petroleum [22, The Schlumberger Organization]. Geco-Prakla is one of seven companies within Oilfield Services [23] and is a member of the Reservoir Evaluation product group. The Reservoir Evaluation product group specialises in wireline and seismic services and is the group concerned with the research described herein.

As is explained in [24] and [25] seismic surveying or exploration involves the use of sound waves propagating through the ground to image geological formations. The sound waves are generated either by hammering the ground or by setting off an explosive charge in a hole drilled into the ground. The installation of explosive charges are shown in Figures A.2 and A.3.

The sound waves are detected by seismic receivers called geophones, which generate electrical currents in response to ground motion. Careful analysis of the sound waves allows the surveyors to determine whether the waves are direct surface-borne, reflected from a subsurface geological interface or refracted along the top of such an interface. Further analysis will, in the specific case of Schlumberger, allow the surveyors to determine if the area contains exploitable petroleum reserves [24].

Seismic work is labour intensive, requiring men and women to work long hours in difficult terrain and weather conditions. Because of terrain conditions and safety requirements the initial surveying work is often done on foot, as are the deployment and retrieval of the

Figure A.1: A hazard encountered by a Schlumberger field worker (Photo: Ryan Chladny; Rocky Mountain House; 1999).

Figure A.2: Seismic explosives drilling rig (Pincher Creek, AB; 1999).

Figure A.3: Sealing seismic explosives behind rig (Pincher Creek, AB; 1999).

geophones and associated cables.

A.1.1 The Navpac

Before seismic measurements can be made a survey of the terrain is performed. Schlumberger has developed, in-house, an inertial navigation system to obtain highly accurate measurements during surveying. The Navpac[1] is a "... portable, strapdown [system] that enables positioning in environments where the global positioning system (GPS) is not effective and where conventional surveys are particularly difficult" [26].

The Navpac uses inertial measurements to calculate the distance and direction of travel of the surveyor wearing it. At set intervals the surveyor will re-calibrate the Navpac's inertial navigation sensor, as shown in Figure A.4. After an afternoon's worth of surveying, the measurements are accurate to within a few meters.

The author and his colleagues visited Schlumberger's operations outside Pincher Creek, Alberta in the summer of 1999. While there the use of the Navpac by Schlumberger surveyors was observed. Although a remarkable tool it suffers from at least one important shortcoming: the person wearing the Navpac must remain extremely still during the relatively frequent system recalibrations. This is difficult when the operator is on sloping

[1] Navpac is a mark of Schlumberger

Figure A.4: Allan Chatenay calibrating a Navpac (Pincher Creek, AB; 1999).

terrain or, as is the case near Pincher Creek, being buffeted by heavy winds. As shown in Figure A.4 Navpac operators sometimes resort to sitting to achieve the required stillness. If the Navpac could be mounted on a stable mobile platform such as the Polar Bear robot it would be less susceptible to environmental conditions which would otherwise hinder its calibration.

A.1.2 Seismic Cable

The deployment and retrieval of seismic cable is extremely important to the seismic surveying industry. The geophones attached to long lengths of interconnected cables are used to read the vibrations used in surveying work.

The distribution, deployment and collection of the seismic cable is labour intensive and time consuming. This is one of the major bottlenecks in seismic work which could be improved with a degree of clever automation.

In forested areas seismic cable is deployed along cutlines, as can be seen in Figures A.5 and A.6. Cable is also often deployed in farmers' fields, along the side of roads and

Figure A.5: A typical narrow "cutline" trail (Photo: Ryan Chladny; Rocky Mountain House, AB; 1999)

Figure A.6: A typical wide "cutline" trail (Photo: Ryan Chladny; Rocky Mountain House, AB; 1999)

Figure A.7: A seismic cable spooler (Photo: Ryan Chladny; Airdrie, AB; 1999)

elsewhere. Most of these areas are accessible using small off-road vehicles.

In any given operation kilometres of seismic cable are deployed and collected by hand. After collection the cable is untangled, repaired and manually spooled before the next deployment. A typical manual spooler can be seen in Figure A.7.

A.1.3 Land Seismic Vibrators

Schlumberger operates Desert Explorer vibrator vehicles in areas where using explosives for seismic work is not a good choice. These vehicles using hydraulic rams to hammer the ground, produce the sound waves required by the geophones.

The Desert Explorers are often driven in formations with one driver in each vehicle. One vehicle is designated as the lead vehicle while the others follow.

The profitability of this activity could be increased through partial automation. As is discussed in the final segment of [27] by removing drivers from the non-lead vehicles and replacing them with navigational systems capable of tracking the lead vehicle driven by a person, costs could be reduced significantly.

Appendix B

State of the Art in Mobile Robotics

B.1 Outdoor Mobile Robotics

When considering a robotics project the general problem is constrained in three ways:

- Mobility: fixed or mobile,
- Environment: land, air, sea or space,
- Level of Autonomy: tele-operated, semi-autonomous or fully autonomous.

This project focuses primarily on mobile, land-based and semi-autonomous robotic systems. The background material covers more than this subset, as research in other branches of outdoor mobile robotics are relevant to this thesis.

B.1.1 Tesla: An Early Pioneer

Arguably the first robot, Nikola Tesla's "telautomaton", was invented in the late 1890's. Like the other roles Tesla played in the development of modern technology his contribution to the field of robotics has largely been ignored. In 1898 he was granted US Patent No. 613,809, "Method of and Apparatus for Controlling Mechanism of Moving Vessels or Vehicles". The telautomaton, a radio-controlled boat, was demonstrated to the public at New York City's Madison Square Garden that same year.

Two things are important about Tesla's work. First, his work has largely been ignored by both the general public and the robotics community. His work has never been mentioned in any course, presentation, magazine, book, radio or television program, or on any web site related to the field of robotics that the author has come across.

The second important fact is that this pioneering project took the form of an outdoor mobile robot. Unfortunately, the work did not progress much further than the original demonstration for at least two reasons. First, there was a lack of interest by Tesla's american investors in using the telautomaton as a weapon in the Spanish-American War. Second, a fire gutted Tesla's New York laboratory, seriously disrupting many of his ongoing projects, including the more important work on electrical power and radio.

B.1.2 "Congress Has Called For a Horde of Unmanned Systems"

In spite of an initial reluctance to accept roboticized systems, the United States military is presently one of the most important sources of funding for the research and development, and deployment of robotic vehicles in the world. A look in copies of the Armed Forces Journal International [28] reveals an immediate military demand for aerial robotic vehicles. Articles on politics behind military funding such as "Out of the Loop – [The United States] Congress Has Called For a Horde of Unmanned Systems..." [28, page 14] are surrounded by military product advertisements which catch the readers' attention with phrases like "There is no section titled 'The Unfair Use of Technology' In the Geneva Convention" and "The Threat Stops HERE."

In addition to the innumerable cruise or guided missile aerial robotic weapon systems, there are a growing number of robots known as unmanned aerial vehicles (UAV) used for reconnaissance and even combat.

> "The first surveillance UAVs were developed by America after the shooting down of Gary Powers's U-2 spyplane in 1960, and they have since played big roles in the Gulf war and in last year's conflict in Kosovo. France, Israel and Britain are also keen on UAVs" [29].

Systems like Aeronautical Systems' Predator and I-Gnat aerial reconnaissance vehicles have been used in North Atlantic Treaty Organization (NATO) actions in Kosovo and other "combat area deployments" [28, page 17]. An important recent development is the Boeing X-45A, officially unveiled in September 2000, the first of a new breed of aircraft: the unmanned combat aerial vehicle (UCAV). As military and political decision makers become more familiar with the aerial robotic technology it is relatively safe to assume that growth

in this sector will continue. In fact, one defence consultant predicts that up to 90% of aerial vehicles will be unmanned by 2025 [29].

Until recently outdoor robotics work has largely been confined to automating aerial and underwater [30] [31] [32] vehicles. The three-dimensional environments in which these robots operate are largely free of obstacles, making their control (if you ignore the complexities of, for example, fluid dynamics) relatively easy.

B.1.3 Ground Robotic Systems

Today, many groups are actively researching technology related to automating outdoor ground mobile navigation. Three major application branches have developed: commercial/roadway, military, and space exploration. In North America and in particular, the United States, these branches share many of the same funding and research resources. In the United States and Canada there are many agencies and offices involved in mobile robotic technology including, but not limited to:

- United States
    - Unmanned Ground Vehicles/Systems Joint Project Office (UGV/S JPO)
        * Operates from the Pentagon in Washington, D.C.
        * Coordinates funding for various projects of military significance [33]
    - Defense Advanced Research Projects Agency (DARPA)
        * Funds academic and non-academic research work of military significance
    - National Science Foundation
        * Funds robotics research, including joint projects with DARPA such as Tactical Mobile Robot [34]
    - Department of Energy (DOE)
        * "Unfocused"[1] robotics research for nuclear and defence industries
    - National Aeronautics and Space Administration (NASA)

[1]Description given by a Sandia National Laboratories administrator during a conversation in Orlando, Florida, USA in July 2000.

      * Space Telerobotics Program

- Canada
  - Institute for Robotics and Intelligent Systems (IRIS) & PRE-Competitive Advanced Research Network (PRECARN)
    * Funds academic and non-academic robotics research
  - Canadian Space Agency
    * Involved in robotics projects such as MacDonald Dettwiler's Space Station Remote Manipulator System
  - Defence Research & Development Canada (Department of National Defence)
    * Research for ground mobile robotic systems, especially for land-mine detection.
    * Focus is primarily on system integration but also includes perception, control, platforms, human-machine interface and communication.[2]
    * Work is primarily conducted at Defence Research Establishment Suffield (DRES)

Practical developments in automating highway travel began in the mid 1980's with the establishment of United States Defense Advanced Research Projects Agency's (DARPA) Autonomous Land Vehicle Project. DARPA has funded such projects as the Transitional Unmanned Ground Vehicle (TUGV) and the Demo I, II and III eXperimental Unmanned Vehicles (XUV).

Carnegie Mellon University (CMU) began collaborating with DARPA in the mid 1980's on the Navigational Laboratory (NavLab) project: a modified General Motors Corporation Chevrolet van capable of driving autonomously under highway conditions. NavLab research has kept pace with developments in computer and sensor miniaturisation and military demand, resulting in it being used on board High Mobility Multi-purpose Wheeled Vehicles (HMMVW, better known as a "Hummer" or "Humvee") and recently, a small commercial sedan.

---

[2]From a discussion with Chris Brosinsky at SPIE AeroSense Conference, 2000

Development of outdoor mobile robotics outside North America is also occurring. In fact, so much Asian and European development in all robotic sub-fields is occurring that on November 5, 1997 the United States Senate Task Force on Manufacturing released a letter urging some of the agencies listed above to work together to "enable the United States to regain its dominant position in the robotics and intelligent machines industry" [35].

One of the german initiatives was described at the AeroSense 2000 conference held in Orlando, Florida, USA in April, 2000 [36] and [37]. Stephan Baton of DaimlerChrysler AG's Dornier division discussed the Program of Intelligent Mobile Unmanned Systems (PRIMUS) robotic project, a modified Wiesel 2 armoured vehicle. Using LASER-RADAR (LADAR) and traditional camera imaging the system has shown automated road as well as contour driving and waypoint navigation capabilities. A live demonstration of the vehicle driving at speeds of up to 50 km/h has been conducted on roads near Pfollendorf, Germany.

The Unmanned Ground Vehicles/Systems Joint Project Office (UGV/S JPO), formed in 1988 by the US Army and Marine Corps has also been actively developing robotic ground vehicles through projects such as the TUGV, Mobile Detection Assessment and Response System (MDARS) Interior and Exterior (security and inventory control projects), and Demo III XUV.

SARGE (Surveillance And Reconnaissance Ground Equipment) [38] and its predecessor Dixie [39], were developed in the early 1990's by Sandia National Laboratories using funding from the Project Office as tele-operated unmanned vehicles for use in hazardous military situations.

CMU has continued its NavLab project and has also collaborated with DRES on equipping a Humvee with the sensor and computational components required for autonomous robotic work as part of the TUGV project [40]. TUGV research was conducted by various groups in North America until about 1998, including CMU, Sandia National Laboratories in California and New Mexico, the United States Tank Armament Command (TACOM) in Michigan, and Canada's Defence Research Establishment Suffield.

The National Aeronautics and Space Administration (NASA) has also been conducting mobile robotics research through its Space Telerobotics Program. In 1998, a series of Antarctic experiments was conducted in conjunction with French and Italian scientists focusing on the "autonomous search of Antarctic meteorites" [41]. It also provided a chance

to demonstrate advances in perception, and essentially autonomous control and navigation under harsh environmental conditions.

These projects display a large variation in levels of autonomy. The Demo III XUV and the Nomad projects are attempting high levels of semi-autonomy and low levels of complete autonomy within the context of well-defined tasks. MDARS, NavLab and PRIMUS are essentially semi-autonomous systems since they act near-autonomously in relatively well defined environments. Likewise, indoor mobile robots such as RHINO [42, Ch. 1], Xavier [42, Ch. 4] and DERVISH [42, Ch. 3] or MDARS Interior [11, page 22], Modbot [11, page 17] or Robart II [11, page 15] are capable of near autonomous behaviour because their environments are relatively structured and simple. TUGV and SARGE were transitional projects for stakeholders in the robotics industry adapting to the issues of semi-autonomy after the primarily tele-operated projects of the 1970s and 1980s.

In outdoor mobile robotic systems where a high degree of autonomy is required there is often a reliance on a complex and expensive assortment of sensors. For example, NASA's Nomad uses four non-scanning laser range-finders, a scanning laser range-finder, a panoramic camera, a stereo camera, and a specialised "science" camera. A powerful processing back-end is required to integrate and analyse the data gathered by the sensor array. All this, and the number of support personnel required leads to high mission costs.

B.1.4 Semi-autonomy: Augmenting People

The introduction of robotics into any given environment does not necessarily imply the elimination of roles for people. Instead, it can give them new tools to avoid dangerous tasks or to perform difficult ones more effectively.

Chris Brosinsky, the Autonomous Land Systems leader for Defence Research and Development Canada, in a presentation given at AeroSense 2000 illustrated some of the priorities for the Canadian Armed Forces while conducting field operations. The idea he tried to convey is that when conducting dangerous military tasks human lives can be saved by placing robotic vehicles in harm's way instead.

The points illustrated in Mr. Brosinsky's presentation were echoed by another presentation given by Mike Toscano, Coordinator of the Pentagon's Joint Project Office (or Joint Robotics and Physical Security Equipment Programs) [43], in which he pointed out that

the US military's proposed Future Combat System is to be used in "dirty, dangerous, dull and impossible missions" [44].

It is obvious that no military force will abandon soldiers and completely replace them with robots any time soon. Robotic systems are, instead, augmenting the effectiveness of soldiers by playing very specialised roles in either attack (e.g. cruise missiles), reconnaissance (e.g. Demo III XUV ground vehicle, Predator aerial vehicle [45]), or urban warfare (e.g. the DARPA Tactical Mobile Robot program). In all of these cases, people are still required but their roles change to accomodate the introduction of robotic devices.

DARPA's Tactical Mobile Robot (TMR) [34] program focuses on small autonomous and semi-autonomous systems, including Science Applications International Corporation's (SAIC) Small Unit Robot (SuBot). The SuBot is a prototype "ThrowBot", a hand-held robot that can be thrown into dangerous environments through open windows or doorways. The Distributed Robotics program funded by DARPA has yielded a similar robot that can be fired from a grenade launcher [33].

From a civilian point of view an example of the augmentation philosophy can be found in some commercial initiatives such as those outlined in [46]. One example is the development of wheel loaders by Caterpillar Corporation. Used in quarries or in mines, Caterpillar's experimental wheel loaders use a robotic subsystem to improve the efficiency of the excavation process. The subsystem takes over from the operator during a portion of each excavation, reducing slippage of the wheels as the shovel gathers a load. Tests with some of the most experienced operators have shown a significant reduction in wear of loader tires. Because each tire can cost up to $30,000 this translates into a significant potential cost savings.

B.1.5 Polar Bear, the University of Alberta's Robotic Vehicle

The Polar Bear, a hydraulically actuated and gasoline powered robot has been designed and constructed by undergraduate and graduate students at the University of Alberta to both compete in an international robotics competition and to attract commercial interest in mobile robotic applications.

The Polar Bear provides a test bed for rugged terrain mobile robotics research, a field which has been maturing in the past fifteen years. At this stage, most outdoor mobile

robotic systems are built in-house due to a lack of commercially available systems. Component selection for Polar Bear focused on commercially available and readily serviceable components while maintaining low cost.

Appendix C

The Autonomous Robotic Vehicle Project

C.1 Two Group Projects: ARVP and SAGV

This development of the Polar Bear robot was driven by two goals. The original goal was to enter the Polar Bear in the annual Intelligent Ground Vehicle Competition (IGVC), an intercollegiate engineering competition held in the United States which promotes the development of outdoor mobile robots. This was the primary goal for the majority of the ARVP student volunteers working on the Polar Bear.

The second goal, which paralleled the first, was to conduct practical robotics research for Schlumberger Geco-Prakla based on the work by the ARVP volunteers. The students who undertook the SAGV research were able to work full-time on the development of the Polar Bear, while also contributing to the ARVP goals. The ARVP goals, with earlier deadlines due to the IGVC being held in early summer, provided timely testing and development milestones for the SAGV research.

C.1.1 Cutting Our Teeth

The development of the Polar Bear provided a huge personal reward to those who showed the initiative to become involved. What started as a solely volunteer project led to the development of new skill sets, employment and research opportunities as well as media recognition.

Many of us learned how to release the magic blue smoke found inside circuitry, while others discovered how to leave fingerprints on freshly welded steel. We all learned the

difficulties of working as a team. Everyone encountered and overcame the difficulties of combining mechanical components with electrical ones and the difficulty of writing software to control it all.

C.2 The Benefits of Simulation

Much of the work done on the project was accomplished without the aid of outside expertise or tools. Some software packages were used and, of these, Pro/Engineer (mechanical engineering), gcc (software development) and Protel (electrical engineering) were the most popular.

To be honest, simulation tools have largely been ignored in the design of the actuator and sensor electronics as well as the software for this project.

An article [47] on the design of automotive electronics discusses statistics regarding common sources of error. The article refers to a speech given by BWM's Dr. Wolfgang Reitzle, who noted that the cost of today's electronics has increased due to the following factors:

- Conceptual errors: 51%
- Hardware errors: 22%
- Logical mistakes: 14%
- Statistical errors: 13%

The article, based partly on this information, concludes that "more computer-aided simulation is therefore required earlier in the design cycle to reduce these errors and the cost they add to the system."

From the perspective of the electrical and software development, available tools were not exploited enough. Matlab and its extensive control, fuzzy-logic and neural network libraries should have been more closely studied. Electronics Workbench (EWB), a simulation tool for circuitry, was dismissed in favour of Protel's printed circuit board layout functions. In retrospect, usage of EWB prior to or rather than Protel would probably have revealed early conceptual errors in motor driver design and sensor integration.

Appendix D

Illustrations and Photos

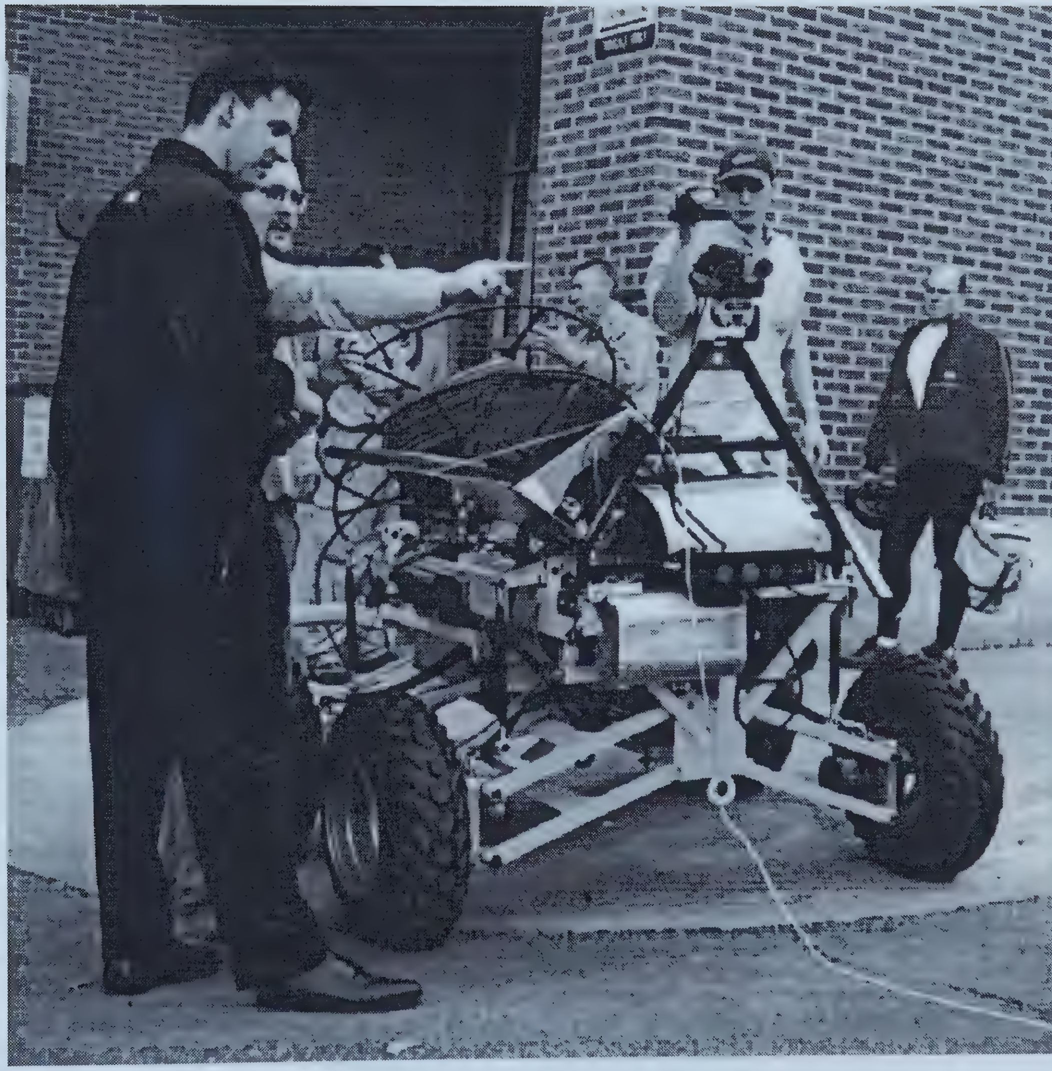

Figure D.1: A frontal view of the Polar Bear during the September 25 2000 demonstration. Allan Chatenay and Kit Barton are on the left, Ryan Chladny is videotaping behind the robot. The Polar Bear's video camera and SONAR array can be clearly seen.

Figure D.2: Jim Qualie wearing the Pursuit Pack and a safety vest.

Figure D.3: Jim Qualie discussing the Polar Bear robot with Allan Chatenay. The Polar Bear's video camera and on-board computer can be seen clearly.

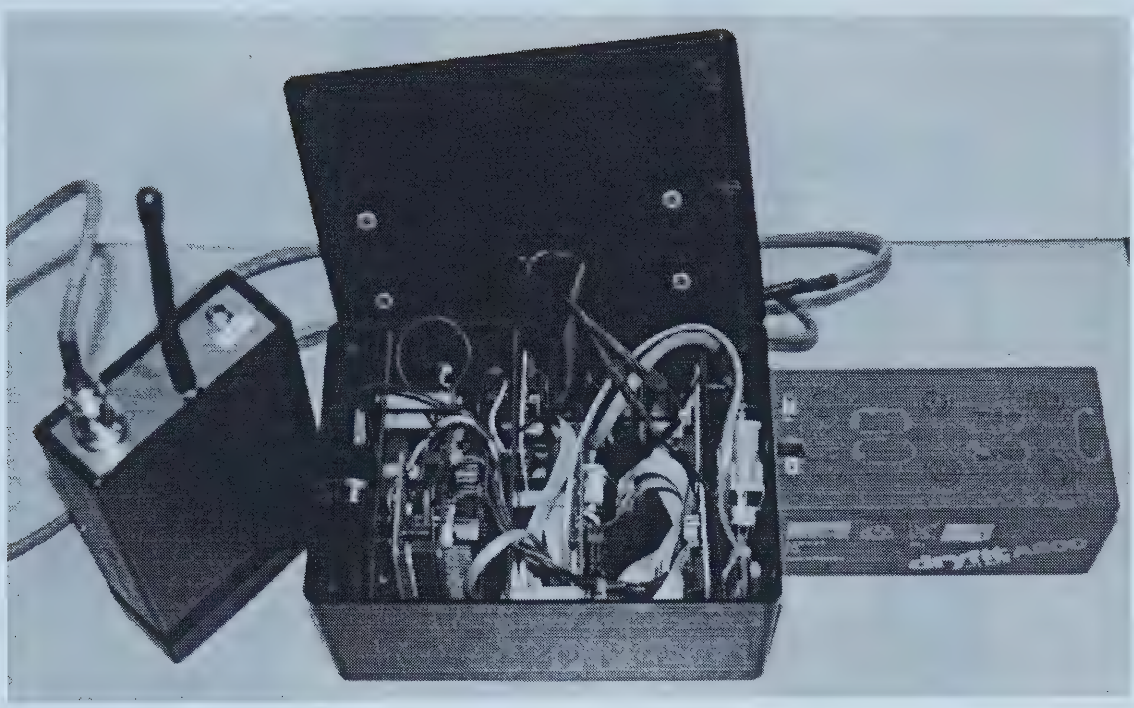

Figure D.4: The Pursuit Pack removed from fanny pack.

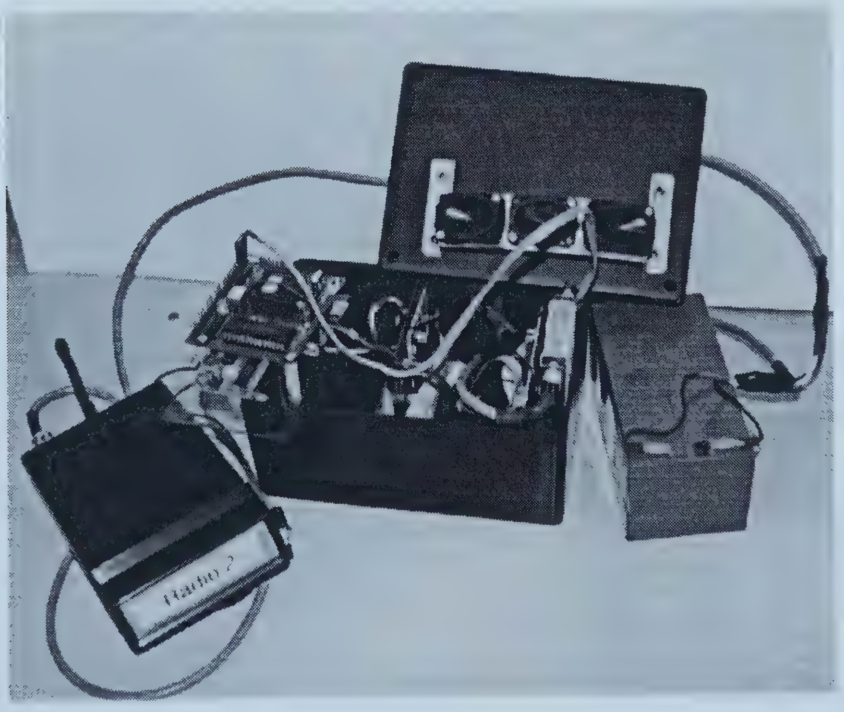

Figure D.5: The Pursuit Pack Components. Micro-controller removed from box.

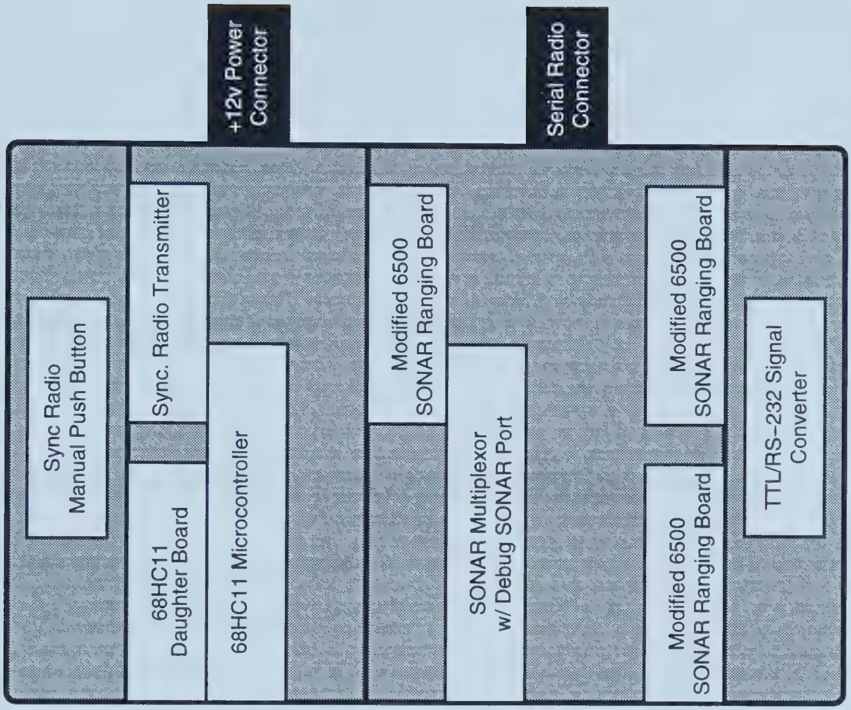

Figure D.6: Layout of subcomponents in Pursuit Pack

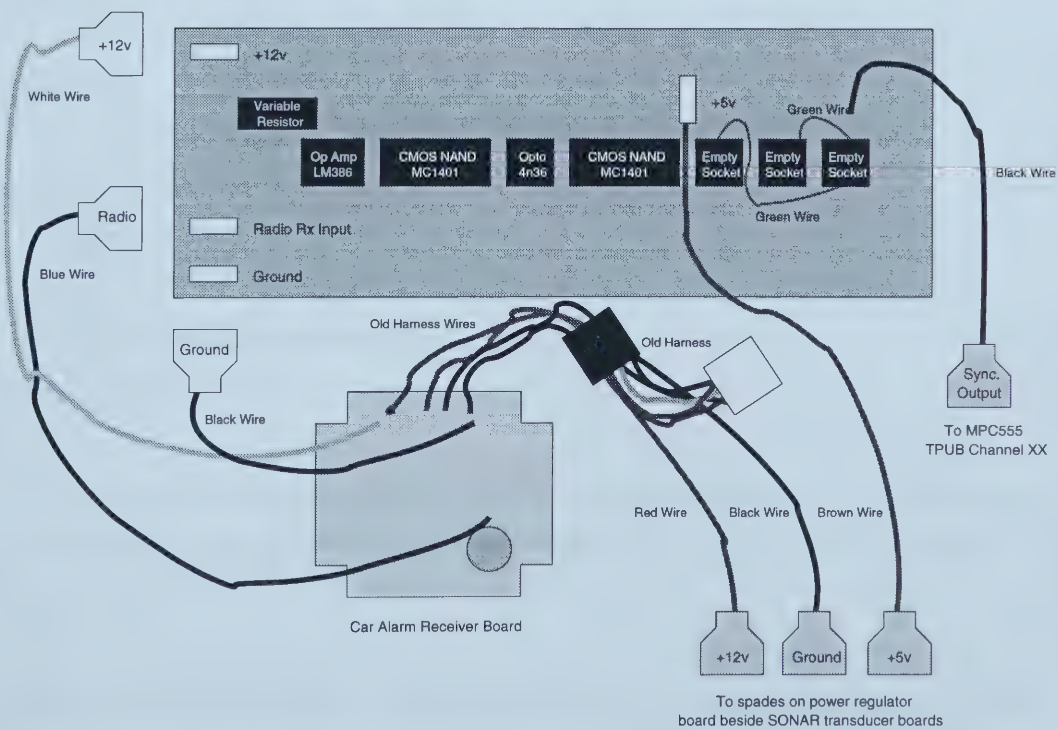

Figure D.7: The Synchronization Radio receiver circuit layout.

Appendix E

Copyright Permissions

E.1 Carnegie Mellon University Millibot Photo

Date: Wed, 11 Oct 2000 03:21:26 -0400

From: Pradeep Khosla pkk@ece.cmu.edu

To: jasmith@ee.ualberta.ca

Cc: rjg@andrew.cmu.edu

Subject: Re: Millibot photos

James

It sounds OK with me. I am cc'ing Bob on this email and he can send you better pictures. He will also send you pointers the proper papers to cite. Thanks

pradeep

At 01:58 PM 10/10/2000 -0600, you wrote:

Hello, Dr. Khosla,

I am writing a thesis on a mobile robot system which uses sonar for localisation. I would like your permission to include some information about the Millibot project in my thesis.

In particular I am interested in discussing the sonar units that are used on the Millibots. Do you have a larger version of the following picture?

http://www.contrib.andrew.cmu.edu/usr/rjg/millibots/mod2.jpg

or

http://www.contrib.andrew.cmu.edu/usr/rjg/millibots/w_sonar1.jpg

Thank you.

James Smith

E.2 Ryan Chaldny's Photos

Date: Fri, 19 Jan 2001 13:15:25 -0700 From: rchladny@maildrop.srv.ualberta.ca To: jasmith@ee.ualberta.ca Subject: RE: Permission

[The following text is in the "ISO-8859-1" character set.] [Your display is set for the "US-ASCII" character set.] [Some characters may be displayed incorrectly.]

I Ryan Chladny, of sound body and mind while under my own volition, give James Smith my complete expressed explicit consent for you to reproduce any and all photos that have been taken by myself relating to his research at the University of Alberta. In exchange for absolutely no cost, monetary or otherwise implied.

Hope that does it you ;)

Good luck again,

RC